# PREFERENCES FOR EFFORT AND THEIR APPLICATIONS 

by
COLIN HARPER CORBETT

## A DISSERTATION

Presented to the Department of Economics and the Graduate School of the University of Oregon in partial fulfillment of the requirements
for the degree of Doctor of Philosophy

June 2016

# DISSERTATION APPROVAL PAGE 

Student: Colin Harper Corbett

Title: Preferences for Effort and Their Applications
This dissertation has been accepted and approved in partial fulfillment of the requirements for the Doctor of Philosophy degree in the Department of Economics by:

William Harbaugh
Van Kolpin
Michael Kuhn
Sanjay Srivastava
and

Scott L. Pratt

Original approval signatures are on file with the University of Oregon Graduate School.

Degree awarded June 2016
(c) 2016 Colin Harper Corbett

# DISSERTATION ABSTRACT 

Colin Harper Corbett
Doctor of Philosophy
Department of Economics
June 2016

Title: Preferences for Effort and Their Applications

In this dissertation, we experimentally examine individual preferences of effort, including time and risk preferences. In Chapter 3, we find that at least in certain settings and mindsets, individuals are very patient in their time preferences for effort, choosing to distribute effort evenly over time periods. However, they do not always live up to the stated plans, suggesting dynamic inconsistency or possibly two separate decision-making systems in the mind. This relates to our model in Chapter 2: a dual-self model of allocating effort between time periods in working toward a larger goal including incomplete information between different mindsets in the same person. Chapter 4 examines the risk preferences for effort, as a measurement of the utility function of effort, and finds that in this setting, subjects are very risk-averse over effort, compared to their preferences over money: they greatly avoid the possibility of having to complete a large number of tasks. These experiments and model help provide an understanding of how individuals allocate the scarce resource of time and energy to tasks they must complete.

## CURRICULUM VITAE

NAME OF AUTHOR: Colin Harper Corbett
GRADUATE AND UNDERGRADUATE SCHOOLS ATTENDED:

University of Oregon, Eugene, OR
University of Rochester, Rochester, NY

## DEGREES AWARDED:

Doctor of Philosophy in Economics, 2016, University of Oregon
Master of Science in Economics, 2013, University of Oregon
Bachelor of Arts in Music and Economics, 2010, University of Rochester Bachelor of Science in Applied Math, 2010, University of Rochester

## AREAS OF SPECIAL INTEREST:

Experimental Economics
Behavioral Economics

## GRANTS, AWARDS AND HONORS:

Graduating Teaching Fellowship, 2011-2016
Dale Underwood Outstanding Graduate Student Scholarship, 2012
Cum Laude, University of Rochester, 2010

## ACKNOWLEDGEMENTS

I thank Professors Harbaugh and Kuhn for their guidance with the experiments and analysis in chapters 3 and 4, and Professor Kolpin with his guidance through the theory in chapter 2. I thank Ben Fitch-Fleischmann for allowing me to run the experiment in chapter 3 in his class, and the University of Oregon Department of Economics for funding the experiment in chapter 4. I also thank my parents and many friends for their support and encouragement throughout my time here.

For all my teachers, in all aspects of life: parents, family, friends, educators, and mentors.

## TABLE OF CONTENTS

Chapter Page
I. INTRODUCTION ..... 1
II. GOALS, REFERENCE POINTS, AND DISAPPOINTMENT COSTS: A MODEL OF SELF-CONTROL ..... 3
2.1. Introduction ..... 3
2.2. Literature ..... 4
2.3. The Model ..... 8
2.4. Possible Outcomes of the Model ..... 26
2.5. Conclusion and Future Work ..... 27
III. ESTIMATING TIME PREFERENCES FOR SCHOOL WORK USING CONVEX TIME BUDGETS ..... 29
3.1. Introduction ..... 29
3.2. Literature Review ..... 31
3.3. Experimental Procedure ..... 36
3.4. Empirical Strategy ..... 41
3.5. Descriptive and Aggregate Results ..... 47
3.6. Discussion and Conclusion ..... 63
IV. RISK PREFERENCES FOR REAL EFFORT TASKS ..... 67
4.1. Introduction ..... 67
4.2. Literature Review ..... 68
4.3. Experimental Design ..... 71
4.4. Test Strategy ..... 78
4.5. Aggregate Results ..... 84
Chapter Page
4.6. Individual Results ..... 96
4.7. Conclusion ..... 99
REFERENCES CITED ..... 101

## LIST OF FIGURES

Figure Page

1. Commitments (missing Family and Social time) ..... 39
2. Allocation Decision ..... 40
3. Allocation within Budgets ..... 48
4. Review Question Allocation ..... 49
5. Real Effort Task ..... 73
6. Effort Comparisons ..... 76
7. Money Comparisons ..... 77
8. Individual Histograms ..... 97
9. Individual Scatter Plots ..... 98

## LIST OF TABLES

Table Page

1. Reduced Form ..... 52
2. Sample Restrictions ..... 54
3. Demographics ..... 56
4. Behaviors ..... 57
5. Day of the Week ..... 58
6. Question Statistics ..... 59
7. NLS Structural Results ..... 61
8. Losses Perspective ..... 64
9. Demographics ..... 78
10. Error Structure Assumptions ..... 85
11. Parameter Estimates for Money ..... 86
12. Parameter Estimates for Leisure ..... 88
13. Parameter Estimates for Effort ..... 89
14. Separated Regressions, Not Reframed ..... 90
15. Order Effects ..... 91
16. Demographic Covariates ..... 92
17. Risk Attitude Covariates ..... 94
18. Personality Covariates ..... 95
19. Individual Estimates ..... 96
20. Individual Correlations ..... 98

## CHAPTER I

## INTRODUCTION

Economics, the study of the allocation of scarce resources, analyzes how individuals allocate their personal scarce resources and their preferences over these resources. However, most analysis of personal scarce resources has been about money: while money is a convenient tool and very important for many economic decisions, people make important economic decisions unrelated to money throughout their daily lives. One primary domain of this is with time and effort: individuals choose how to allocate the scarce resource of their time and physical and mental energy on completing tasks and working toward goals. In short, they are allocating effort. This dissertation aims to analyze how individuals make decisions about effort on a personal level, including their preferences for effort and decision-making processes behind it.

In Chapter 2, Goals, Reference Points, and Disappointment Costs: A Model of Self-Control, we create a model to analyze the behavior of individuals working toward a long-term goal, using a dual-self formulation in a game-theoretic setting. Two decision-making processes in the same individual persons mind compete: a long-term planning self, and short-term acting self. What separates this model from other dual-self models is the incomplete information between the two selves: neither side is fully informed about the time preferences of the other side, and thus cannot fully predict actions the other player takes. This can lead to failing to live up to plans, over-shooting when making plans, and other seemingly irrational behavior. Elements of this model including time discounting of the disutility of effort and risk preferences (utility curvature) of effort, both of which are not well-studied in
the literature. The next two chapters analyze those two elements of the individual decision-making processes behind effort.

Chapter 3, Estimating Time Preferences for School Work Using Convex Time Budgets, analyzes time preferences for a specific type of effort: school work in the form of multiple choice review questions. In a field experiment within an introductory economics class, we analyze how students make decisions about allocating school work over time. We find that students are surprisingly patient when allocating tasks: they show a strong preferences for evenly allocating tasks between time periods, especially when they complete the experiment early within the survey period. However, when actually following through with these plans, some students do not complete their assigned tasks. This provides some evidence of the dual-self model of decision-making: students in a long-term mindset choose to allocate effort evenly over time, then the short-term mindset fails to complete those plans, aligning with predictions from the model in Chapter 2.

We analyze another major element of individual preferences in Chapter 4: Risk Preferences for Real-Effort Tasks. In a lab experiment with two sessions, subjects compared lotteries over gains and losses of money (in one session) and effort/leisure (in the other). We find that subjects are considerably more riskaverse over leisure (the absence of effort) than money, and that there was only small correlation between risk preferences for money and those for leisure. Knowing the shape of the disutility function of effort, especially in contrast to that of money, allows us to make better predictions about individual behavior in settings like the Chapter 2 model.

## CHAPTER II

## GOALS, REFERENCE POINTS, AND DISAPPOINTMENT COSTS: A MODEL OF SELF-CONTROL

Many models of self-control exist and make various predictions about the consequences of dynamic inconsistencies, effort provision, and reference-based utility. In this paper, we combine many of these threads into a model of longterm goal setting and effort provision. Built on a framework of short-term actions affecting long-term states, this model stylizes goal-setting and self-control as a game played between two states of thinking, the long-term and short-term selves. Underlying this model is a theory of reference-based utility of actions using goals as reference points, labeled as a disappointment cost. Using aspects of this model, we are able to predict and rationalize apparent failures of self-control. We provide a simplified example of procrastination while writing a research paper.

### 2.1 Introduction

People set long-term goals and make plans to complete these goals all the time, and sometimes complete them, but sometimes don't. They may procrastinate along the way, choose commitment devices to assist them, or work to avoid falling short of their expectations for themselves. Sometimes these goals are completed, but sometimes they are not. What explains how people fail to live up to their own plans of action? Naively time-inconsistent preferences, high costs of self-control, loss aversion, and more have been proposed as reasons for generating this seemingly irrational behavior of unmet goals. This paper attempts to create a model of goalsetting, action-planning, and execution that combines many of these elements that
shows how agents, even when rationally setting strategies to maximize utility, can be left with unmet goals and disappointment. Section 2 reviews the various threads of the literature contributing to this research, section 3 proposes and defines all aspects of the model, section 4 discusses some of the possible outcomes of the model resulting from various internal assumptions, and section 5 concludes.

### 2.2 Literature

This model combines three ongoing threads in the behavioral economics literature: the study of self-control and commitment, utility of effort and leisure, and reference utility, and attempts to combine them into a single model of goal-setting and effort provision. Just like many other models, we include time preferences with possible dynamic inconsistencies, another topic of extensive research.

## Self-Control and Commitment

Many previous authors have created theories about long-term decision-making regarding control and commitment. This includes models of procrastination, which is often a failure to control oneself. One example of procrastination that does not involve a control failure is Fischer (2001), which studies procrastination under timeconsistent preferences. However, time-consistent preferences don't induce dynamic inconsistencies, which the term "procrastination" usually implies. A more familiar theory of procrastination comes from O'Donoghue and Rabin (2001), who define a model of procrastination, and extend it in (2008). Here they introduce the idea of agents having incomplete awareness over their own dynamic inconsistency (as opposed to complete awareness or complete naiveté). We model this incomplete self-awareness as unknown player types in a Bayesian game.

With some awareness of their own dynamic inconsistency, agents would be expected to desire some sort of commitment device or otherwise struggle for selfcontrol - various ways agents counteract dynamic inconsistencies. Laibson (1997) creates a theoretical commitment device that agents use to to control themselves in future periods, knowing their own hyperbolic discounting, and models these decisions as a game between different selves in different time periods. Gul and Pesendorfer (2001) discuss preferences for commitment devices to counteract temptation or tendencies to procrastinate. Commitment devices, reviewed by Bryan et al. (2010), have been shown to work in reducing dynamically inconsistent behavior in a number of contexts, but in an experiment studying effort provision, Bisin and Hyndman (2014a) find little evidence for effectiveness of external commitment devices to prevent procrastination on real-effort tasks. In our model, we essentially use an internal commitment device in the form of internal disappointment to model behavior when external commitment devices may be unfeasible.

To explain the occurence of dynamic inconsistencies and seemingly irrational behavior, Fudenberg and Levine (2006) and Fudenberg and Levine (2011) introduce a model of dual-selves: a long-term self who makes plans, and a short-term self who acts on these plans. This idea of two seperate decision-making systems in the brain in based on neuroscience evidence (see McClure et al. (2004)), and aligns fairly well with introspection. We borrow heavily from this set-up, but change a number of important assumptions, especially including not assuming the possibility of perfect self-control.

## Reference Utility

The idea of utility being based on reference points first appeared in the idea of prospect theory from Kahneman and Tversky (1979) and Tversky and Kahneman (1991), which addressed the concept of gain-loss utility, and has appeared across various areas of the literature since. However, the topic of how agents formed reference points was left unsettled, and remains an open question. One original idea for how agents form expectations was the status quo, advocated by Samuelson and Zeckhauser (1988). But currently, a leading candidate for a model of reference points is expectations - instead of basing your reference point for today on where one was yesterday, one uses yesterday's beliefs about today as a reference point. Many times these are identical, but often are not. Kőszegi and Rabin (2006) first modeled reference-based utility using expectations as a reference point, and continued on that subject in Koszegi and Rabin (2007) and Kőszegi and Rabin (2009). Matthey (2008) introduces reference-based utility over expectations, which relates to our consistency cost function.

Providing experimental evidence in a controlled setting for expectations as reference points is Ericson and Fuster (2011); others have found similar results in applied microeconomics. Gill and Prowse (2012) find further evidence of choiceacclimated expectations as reference points in a tournament setting, labeling this phenomenon as disappointment aversion. But actions taken by oneself as a reference point have not yet been studied.

## Utility of Effort

Much of the behavioral literature focuses on utility over consumption generally, or money specifically. However, the model proposed in this paper focuses
on effort provision toward a goal, rather than consuming material goods or money.
In one sense, effort provision is loss of leisure, subject to framing effects. But disutility of effort is felt as more than just a loss of leisure: Kurzban et al. (2013) propose that effort provision is costly because executive function is a limited resource subject to opportunity costs. However, the disutility of effort is not directly comparable to consumption utility, so the creation of new models for effort is warranted. From Augenblick et al. (2015) we also know that time preferences for effort-related tasks can be different than time preferences for money and consumption, giving a basis for creating different models to describe effort than consumption.

Abeler et al. (2011) provide some experimental evidence for changing provision of effort based on changing expectation-based reference points, although these reference points are experimentally controlled, rather then endogenously generated.

## Game Theory

While the game theory structure of our model is mostly very standard, it employs a few specific elements worth citing. The structure of the entire model is based on Stackelberg competition based on Stackelberg (published 1934, finally translated 2011). However, because there is uncertainty in player types, we have a Bayesian game, whose structure was standardized by Harsanyi (1967). Stackelberg Bayesian games have received the most interest recently in the field of security, with a defender choosing a strategy to counter an attacker of unknown type. Future work will delve more carefully into this literature for possible applications.

### 2.3 The Model

In this paper, we present a model of interaction between the shortterm perspective of an individual, referred to as the "short-term self", and the perspective more focused on the long term, known as the "long-term self". Generally described, there is a long-term goal over a span of time. In each period, the long-term self selects a plan of action to work toward that goal, then the short-term self sees the plan and chooses an action to take. This is similar to a Stackelberg competition model, with the short-term self taking the long-term self's plan as a given while deciding which action to actually perform, while the longterm self has to take into account the short-term self's possible reactions when forming a plan.

Preliminarily, time periods in this model are most easily imagined as days, but could be extended to weeks or months or even shortened to hours; examples will assume days.

## States and Transitions

The basic structure behind our model is as follows: an individual experiences a state variable at a particular time, $s_{t} \in S_{t}$. Examples include current level of fitness, level of skill on a musical instrument, or personal weight; throughout this paper we will follow the example of progress on a research project. For model tractability, we assume that the state is real-valued, so progress is easily measurable; this could be time to run a mile, percentage of a musical piece mastered, or for our running example of a research project, number of hours invested into the project. This state is affected by actions the individual takes: exercising, practicing that instrument, or working on research. ("Effort" can also
describe the action taken when applicable.) The action taken in time period $t$ is described as $a_{t} \in A_{t}$, and the process through which actions affect the state is described by $h\left(a_{t} \mid s_{t-1}\right)$ - how actions affect the state depends on the current level of the state. While actions may not be real-valued, the way they affect the state must be. Therefore, the state after period t is described as:

$$
\begin{equation*}
s_{t}=g\left(s_{t-1}\right)+h\left(a_{t} \mid s_{t-1}\right)+\epsilon_{t} \tag{2.1}
\end{equation*}
$$

where $g\left(s_{t-1}\right)$ is how the state evolves by itself: often static (level of work completed) or tending toward zero (level of fitness); and where $\epsilon_{t}$ is external variation in the state. This can include both external variation in the state (like computer crashes) and variation in the effectiveness of effort (like a problem being harder than anticipated, or external distractions). Let $q\left(\epsilon_{t}\right)$ be a probability density function for the distribution of possible values of $\epsilon_{t}$. This distribution may vary with time, but for simplicity, we assume that the distribution is independent of all other elements of the model. Because of this stochasticity, agents can rarely make perfect predictions about the evolution of the state, but instead predict distributions. Note that we assume both $g(\bullet)$ and $h(\bullet)$ are time-invariant, that is, the evolution of the state and the effect of actions are the same in all periods. However, it is possible for the effect of effort to depend on the current level of the state when there are diminishing marginal returns of effort, such as in improving fitness levels. We assume there exists a null action $a_{t}^{0}$ such that $h\left(a_{t}^{0} \mid s_{t-1}\right)=0 \forall s_{t-1}$ - the effects of not working at all do not depend on the current state.

The time formulation of equation 1 implies that at the beginning of period $t$, agents see state $s_{t-1}$, and only after the action in period $t$ has been realized does the state $s_{t}$ actually occur. This implies that the relevant state we assign to "today
" occurs at the end of the day. For example, in completing a research project, every day our researcher wakes up, sees the amount of work completed by the previous night $s_{t-1}$, completes a certain amount of work $a_{t}$, then after completing his work, evaluates his progress for that day $s_{t}$. We also assume that the model ends at time period $T$, so $s_{T}$ is the final realization of the state at the end of time period T , which is the final result of interest to the agents.

## Structure and Information

As stated before, our model includes two agents, both aspects of the same individual, the short-term self and the long-term self. We assume that the longterm self is working toward an external goal with a particular deadline, $s_{T}^{*}$ : reaching some level of fitness for an upcoming race, being able to perform a musical piece in a particular performance, or completing a research project by a deadline. Note that this goal is chosen externally from the model, and that the workings of the model do not affect the goal. Although he is working toward a goal, the actual objective of each agent is to maximize expected discounted utility - we assume this utility includes the benefits of reaching the goal and the costs of getting there. While the utility of the long-term self is dependent on reaching the goal, only the short-term self is capable of taking actions $a_{t}$ and thus influencing the state through the transition function, leading the state toward the final goal. In order to influence the short-term self's choice of action, the long-term self sets a plan of action every period, $a_{t}^{*}$, which acts as a reference point while the short-term self is making decisions.

These two agents interact in a Stackelberg-style game: every period, the longterm self acts first by setting the plan $a_{t}^{*}$ to influence the decision-making of the short-term self. Then the short-term self acts second by choosing the action, $a_{t}$,
which may or may not match the set plan. In our example, a researcher may set a plan to work on a paper for five hours on a day, but could fail to achieve that plan and end up only working for three hours. Note that because plans and actions are just strategies in order to maximize utility, it is possible that they won't lead to the stated external goal, if it is not feasible or the actions required to reach the goal are too costly to be justified.

In this model, the short-term and long-term have time preferences over the length of the game, with which they weigh utilities from different periods. Letting $k$ be the subscript designating the short-run self, and $l$ designating the long-run self (with $i$ indexing agents), $\gamma_{t, k}$ and $\gamma_{t, k}$ represent the discount factors for the time period $t$ units into the future for the two agents, respectively (so, in period 3 , preferences for time period 6 would be $\gamma_{3, i}$ ). For simplicity, let us only consider dynamically consistent preferences: the relative discount factors for two time periods do not depend on the reference time period. With this assumption, agents must behave as exponential discounters, with discount factors determined by a single discount rate $\delta: \gamma t, i=\delta_{i}^{t}$ This allows optimal strategies to be time consistent, and not varying with time.

This assumption about discounting differs from previous models of self-control in which only the long-term self was forward-looking, while the short-term self was assumed to be solely focused on the current period. This model generalizes these assumption to allow the short-term self to

$$
\begin{equation*}
\gamma_{k, t} \leq \gamma_{l, t} \forall t \tag{2.2}
\end{equation*}
$$

Or alternatively, given exponential discounting:

$$
\delta_{k} \leq \delta_{l}
$$

These time preferences present the possibility of an information asymmetry in our model: it is likely that the long-term self is not fully aware of the level of

These beliefs can be modeled through an approach from Harsanyi: each agent assumes that his own preferences are determined by his assigned type, and the preferences of the other agent are determined by the other agent's assigned type. Both types are assumed by both agents to be drawn from common-knowledge distributions of possible types. Note that as long as these assumptions are common knowledge between the agents, it doesn't matter if they are "correct": preferences may in reality be completely deterministic, but if agents believe they are randomized, the effects on the model's predictions are the same as if they were truly random. We define these preferences with types: $\theta_{i}$ is the type of agent $i$, drawn from a set of $N_{i}$ possibilities $\Theta_{i}$. The $\theta_{i}$ type defines time preferences $\left\{\gamma_{i, \tau}\right\}_{\tau=t}^{T}$ : a string of discount factors for all periods. $\Theta_{i}$ must include a type corresponding with the true string of time weights for the respective agents, but also includes a finite number of other possible strings that obey our assumptions about time preferences.

Both agents also assume an initial distribution of possible types for both agents, modeled as $f\left(\Theta_{k} \times \Theta_{l}\right)$, from which types were assumed to be initially drawn. (Let $F(\bullet)$ be the set of possible distributions). This distribution is common knowledge. $F\left(\Theta_{k} \times \Theta_{l}\right)$ only includes distributions where the long-term self is always at least as patient as the short-term self: the probability of a draw where the type of the long-term self is less patient than the type of the short-term self is zero. Each agent obviously knows the realization of his own preference type, but
not that of the other agent; however, knowing his own type, each agent can assume a conditional distribution of the other player's type: for example, $f\left(\Theta_{k} \mid \theta_{l}\right)$ is the long-term self's inferred beliefs about the distribution of types of the short-term self (drawn from $F\left(\Theta_{k} \mid \theta_{l}\right)$ ). As an example, if there were only two possible types of time preferences across both agents (say, impatient and patient), then a long-term self who knows he is impatient would know the short-term self has to be impatient too, as the short-term self must not be more patient. But a patient long-term self would not be certain of the short-term players type, as either a patient or impatient short-term self is possible.

On each information set (and thus in every time period), each agent has beliefs about the other player's type: $\beta_{t, i}$. These beliefs are a distribution of types for the other player in a particular time period. The collected beliefs across every time period and information set for each player are modeled as $\beta_{i}$. These beliefs are likely within the set of possible distributions $\left(\beta_{t, l} \in F\left(\Theta_{k} \mid \theta_{l}\right)\right.$ for the longterm self), but having unjustified, impossible beliefs is still covered by the model. In equilibrium, as agents observe actions from the other player, they can update their beliefs about the other player's type from the initial assumed distribution. $f\left(\Theta_{k} \mid \theta_{l}, \bullet\right)$ denotes the equilibrium beliefs of the long-term self about the shortterm self's type, given a particular information set. In a Weak Perfect Bayesian Equilibrium, these beliefs on information sets that may occur with non-zero probability must be "correct": based on the initial common-knowledge distribution of types.

While time preferences are the likeliest source of information asymmetries in this model, a similar modeling strategy could potentially be used for other elements of the utility function, especially with types over disappointment functions or
consistency costs (explained later). However, we will not attempt to model these in this paper.

## Actions and Strategies

In each period, each player makes one move: first, the long-term self sets a plan of action $a_{t}^{*}$. Then, the short-term self selects an realized action, $a_{t}$, chosen from a set of possible actions for that period $A_{t}$. This set describes the actions that the short-term self is capable of making in a particular time period; it could be the number of hours available to work in a day, the number of miles he is capable of running, etc. $A_{t}$ could also be a non-numeric choice space, like selecting which food to eat. We do assume that $A_{t}$ always includes a null action: $a_{t}^{0} \in A_{t}$. The longterm self's plan of action $a_{t}^{*}$ is chosen from the set of possible plans: $A_{t}^{\prime}$. If only feasible plans are possible, then $A_{t}^{\prime}=A_{t}$, but if unfeasible plans are possible (like working 25 hours in a day), $A_{t} \subset A_{t}^{\prime}$. If $A_{t}$ and $A_{t}^{\prime}$ are non-numeric choice sets, then we assume there exists a metric over these spaces with which agents can compare distance, in order to determine consistency costs and disappointment (more on these in section 3.5).

Coming into each period, the information set of the long-term self includes the full history of the state, and all previous plans and realized actions: $\left\{s_{\tau}\right\}_{\tau=0}^{t-1}$ and $\left\{a_{\tau}, a_{\tau}^{*}\right\}_{\tau=0}^{t-1}$. The short-term self term self also observes the sequence of states and previous plans and actions, but additionally observes the plan for the current period: $a_{t}^{*}$. Along with observing these information sets, both agents can use these observations to update their beliefs about the probability distribution of types for the other agent. For example, if the long-term agent observes actions that would typify a present-biased type from the short-term self, he may increase the probability of present-biased types in his believed distribution. Therefore, with
every information set an agent sees, there is an associated belief distribution for that agent.

The strategies of the long-term self are thus mappings from the product of all previous state, action and plan spaces, to the space of possible plans of action, conditional on the type of the long-term self. For the short-term self, strategies are mappings from the same space plus the current period's plan space map, to the current action space, conditional on the short-term self's type. The initial assumed distribution of types of the two different agents obviously factor into forming a strategy, but they are part of the structure of the game and the same for every information set. Also, while the updated beliefs about the type of the other agent will influence the decision-making, these beliefs are fully determined by the initial distribution and the observed states and actions that the agent uses for updating, including them in the mapping does not bring any new information.

We use $\sigma_{i, t}$ to represent the portion of the strategy in time period $t$-contrary to standard game theory notation, we assume only pure strategies are possible, even though we use $\sigma$, as $s$ already denotes the evolving state.

$$
\begin{align*}
& \sigma_{t, l}: \Pi_{\tau=0}^{t-1}\left(S_{\tau}\right) \times \Pi_{\tau=0}^{t-1}\left(A_{\tau}\right) \times \Pi_{\tau=0}^{t-1}\left(A_{\tau}^{\prime}\right) \times \Theta_{l} \rightarrow A_{t}^{\prime}  \tag{2.3}\\
& \sigma_{t, k}: \Pi_{\tau=0}^{t-1}\left(S_{\tau}\right) \times \Pi_{\tau=0}^{t-1}\left(A_{\tau}\right) \times \Pi_{\tau=0}^{t}\left(A_{\tau}^{\prime}\right) \times \Theta_{k} \rightarrow A_{t} \tag{2.4}
\end{align*}
$$

As a shorthand, we use $\omega_{t-1} \in \Omega_{t-1}$ to designate the entirety of information (states, plans, and actions) from the beginning through the end of time period $t-1$, and thus available at the beginning of $t$. Strategies could then be defined more compactly:

$$
\begin{align*}
& \sigma_{t, l}: \Omega_{t-1} \times \Theta_{l} \rightarrow A_{t}^{\prime}  \tag{2.5}\\
& \sigma_{t, k}: \Omega_{t-1} \times A_{t}^{\prime} \times \Theta_{k} \rightarrow A_{t} \tag{2.6}
\end{align*}
$$

This shows the strategy for each player for actions in time period $t$. For ease of notation, we also define $\omega_{t-1}^{*}$ as the information set $\omega_{t-1}$, plus the plan that the long-term self makes in period $t$. The complete strategy for each player is then the product of these one-period strategies across all future time periods.

At any particular information set, each agent is likely to be unsure about the type of the other agent. Observing the plans or actions in the past of the other player given a particular information set, and inferring what optimal strategies would be in those situations given different player types, each agent can use Bayes' rule to determine the likelihood of different types for the other player. We model these beliefs as $f\left(\Theta_{k} \mid \theta_{l}, \omega_{\tau-1}\right)$ : the long-term self's beliefs about the distribution of short-term player types on a particular information set in period $\tau$; the shortterm self's beliefs would additionally be dependent instead on $\omega_{\tau-1}^{*}$. Therefore, each player has beliefs about the other player's type for all information sets; any equilibrium strategy would necessarily have these beliefs associated with it. This belief structure is consistent with a Weak Perfect Bayesian Equilibrium, assuming optimal strategies.

## Payoffs

We assume that in each period, each agent receives utility from the realizations of the various strategies, and works to maximize his discounted expected utility. The expected utilities of each player, determined by his own
strategies and strategies of the other player (including all possible types), are shows below.

Expected discounted utility of strategies for the short-term and long-term selves:

$$
\begin{align*}
& U_{l}\left(\sigma_{l}, \sigma_{k}, t \mid \omega_{t-1}, \theta_{l}, \beta_{l}\right)= \\
& E_{t, l}\left[\sum_{\tau=t}^{T} \gamma_{\tau-t, l}\left(\theta_{l}\right)\left[m_{l}\left(\sigma_{k, \tau}(\bullet)\right)+n_{l}\left(s_{\tau}(\bullet)\right)+C\left(\sigma_{l, \tau}(\bullet) \mid \sigma_{k, \tau}(\bullet)\right)\right] \mid \beta_{l}, \sigma_{l}, \sigma_{k}\right]  \tag{2.7}\\
& U_{k}\left(\sigma_{k}, \sigma_{l}, t \mid \omega_{t-1}^{*}, \theta_{k}, \beta_{k}\right)= \\
& E_{t, k}\left[\sum_{\tau=t}^{T} \gamma_{\tau-t, k}\left(\theta_{k}\right)\left[m_{k}\left(\sigma_{k, \tau}(\bullet)\right)+n_{k}\left(s_{\tau}(\bullet)\right)+D\left(\sigma_{k, \tau}(\bullet) \mid \sigma_{l, \tau}(\bullet)\right)\right] \mid \beta_{l}, \sigma_{l}, \sigma_{k}\right] \tag{2.8}
\end{align*}
$$

$$
\begin{aligned}
& \text { where } \begin{aligned}
s_{\tau}(\bullet) & =g\left(s_{\tau-1}\right)+h\left(\sigma_{k}(\bullet) \mid s_{\tau-1}\right)+\epsilon_{\tau} \\
\sigma_{l, \tau}(\bullet) & =\sigma_{l, \tau}\left(\omega_{\tau-1}, \theta_{l}\right) \\
\sigma_{k, \tau}(\bullet) & =\sigma_{k, \tau}\left(\omega_{\tau-1}^{*}, \theta_{k}\right)
\end{aligned}
\end{aligned}
$$

Expectations are taken over the other player's type (the long-term self takes expectations over $\theta_{k}$ and the short-term self takes expectations over $\theta_{l}$ ) based on beliefs at each information set, and over the stochastic element $\left(\epsilon_{t}\right)$.

These equations are complicated; we will describe each element of the equation in the following sections.

First, the arguments of the function: utilities for each agent $\left(U_{k}\right.$ and $\left.U_{l}\right)$ are dependent on the strategies of each agent, conditional on the current state, the given player types, and the beliefs each player has about the other's type in that period.

From the current time period $t$, agents look into future time periods, indexed by $\tau$, and discount the utility they expect to receive in those time periods by $\gamma_{\tau-t, i}$ - that is, the relative discount rate between the current time $t$ period and future period $\tau$. The method of forming these expectations is described in section 3.6.

Here, $\sigma_{l, \tau}\left(\omega_{\tau-1}, \theta_{l}\right)$ notates the plan $\alpha_{\tau}^{*}$ made by the long-term self at time $\tau$, given type $\theta_{l}$ and history $\omega_{\tau-1}$. Similarly, $\sigma_{k, \tau}\left(\omega_{\tau-1}^{*}, \theta_{k}\right)$ is the action $\alpha_{\tau}$ taken by the short-term self at time $\tau$, given type $\theta_{k}$ and history $\omega_{\tau-1}^{*}$, which includes the plan made by the long-term self in period $\tau$. Note that histories (and thus information sets) that occur in the future have a random element: $\omega_{\tau-1}$ is unknown to each player, and each player can only take an expectation over possible values of that information set.

The (undiscounted, pre-expectation) contribution to utility from a single period $\tau$ is shown as follows, presented here for simplicity only in terms of actions and plans, rather than strategies.

$$
\begin{gather*}
u_{\tau, l}=m_{l}\left(a_{\tau}\right)+n_{l}\left(s_{\tau}\right)+C\left(a_{\tau}^{*} \mid a_{\tau}\right)  \tag{2.9}\\
u_{\tau, k}=m_{k}\left(a_{\tau}\right)+n_{k}\left(s_{\tau}\right)+D\left(a_{\tau} \mid a_{\tau}^{*}\right) \tag{2.10}
\end{gather*}
$$

Now, let us explain each element in these utility functions. We have already seen how the time preferences of each agent are determined by the type of the player, and the assumptions we place on this string of time preferences: each agent weights each period no less than the next one, and the short-term self weights future periods no more than the long-term self does.

In this utility function, both the short-term self and long-term self derive (dis)utility from both the action taken in that time period and the resulting state that occurs in each period. Utility from actions is notated $m_{k}\left(a_{t}\right)$ and $m_{l}\left(a_{t}\right)$ for the short- and long-term selves, respectively, and utility from states is $n_{k}\left(s_{t}\right)$ and $n_{l}\left(s_{t}\right)$. It is possible for these utility functions to be time-dependent (like working on research being harder on a Monday than a Wednesday), but we will not include that in our notation, for simplicity.

We assume that taking non-null actions (that is, effort) is costly but success is enjoyable: $m_{i}\left(a_{t}^{0}\right)=0$, while $m_{i}\left(a_{t}\right)$ is non-increasing with increasing (positive) distance from the null action (so disutility is non-decreasing), while $n_{i}(\bullet)$ is nondecreasing with increasing progress from zero, toward the goal. Essentially, subjects are less happy having to work harder, but they like making progress. Also, both selves experience increasing marginal costs of effort: in the simple case if $A_{t}$ is one-dimensional and continuous, then $m(\bullet)$ is concave (or the negative of effort costs is convex). In our example, increasing the number of hours of research in one day decreases the utility of that day, with each hour of research becoming more unpleasant.

The long-term goal $s_{T}^{*}$ fits into the model by affecting the state-based utility of the long-term self (and possibly that of the short-term self): the utility of reaching the goal at the deadline is much greater than the utility of not reaching
it: $n_{l}\left(s_{T}^{*}\right) \gg n_{l}\left(s_{T}\right) \forall s_{T}<s_{T}^{*}$ - there is a large discontinuity in $n_{l}(\bullet)$ at $s_{T}^{*}$. This drives the long-term self to maximize his utility by attempting to reach this goal.

## Consistency Costs and Disappointment

The final two unexplained elements of the utility function are $C(\bullet)$ and $D(\bullet)$, known as Consistency Cost and Disappointment, respectively.

We assume that when setting goals, the long-term self takes into account how closely the realized actions of the short-term self will match the set plan. When the short-term self fails to live up to the plan, the long-term self experiences some utility penalty. We model this as a Consistency Cost: $C\left(a_{t}^{*} \mid a_{t}\left(\theta_{k}\right)\right)$. Here, the longterm self compares his stated plan to the action it would induce for all possible types of the short-term self, having a particular consistency cost for each possible short-term self type. (This cost then enters the expectation function as all other terms). Essentially, if the short-term players actions are not consistent with the plan he created, the long-term self experiences a cost. We assume this consistency cost function is zero when $a_{t}^{*}=a_{t}$, and decreasing (so increasing in absolute value) with greater positive distance between action and plan - the farther the action undershoots the plan, the more costly it is. There could likely be a discrete jump in consistency cost as soon as $a_{t}^{*}>a_{t}$, but we don't require it in the model. When the action is greater than the plan, however, this function could be positive, negative, or zero, depending on how the long-term self feels about setting too low of a plan.

Note that this utility cost is not generated by surprise (other than an unlikely type draw), as in equilibrium the long-term self knows what the short-term self's reactions will be for each possible type. There could be many psychological explanations for this cost: one idea is that the long-term self believes his reputation suffers when his plans are not obeyed. Or he feels ineffective and bad about his
inability to influence the short-term self. Another idea is that setting a plan has effects outside of just telling the short-term self what to do, including external enforcement mechanisms, or getting irrationally attached to outcomes from the plan. At the core of this function is that the long-term self uses the shortterm self's action as a reference point from which to evaluate utility of his plan, although the location of this reference point depends on the short-term self's type. Essentially, the effect of this consistency cost is preventing the long-term self from setting unrealistically high plans.

The driving force of this model, however, is the assumption that the shortterm self feels disappointment from failing to achieve the plans set by the long-term self, and this disappointment has a utility cost. Because the short-term self chooses the realized actions, it is thus the short-term self who compares these actions to the plan for the current period, and thus may experience disappointment, depending on this comparison. And it is this feeling of disappointment that drives the short-term self to follow the plans of action created by the long-term self; this feeling serves as a sort of commitment device when no outside devices are possible. Disappointment for period $t$ is modeled by the function $D\left(a_{t} \mid a_{t}^{*}\right)$; it depends on the comparison between the realized action and the planned action; the plan acts as a reference point from which the short-term self effectively feels loss aversion. We assume that $D\left(a_{t} \mid a_{t},{ }^{*}\right)=0$ if $a_{t}=a_{t}^{*}$ : that is, when the short-term self successfully implements the plan of action, he feels no disappointment. This feeling of "disappointment" we are describing is not disappointment with the state of the world, but rather disappointment with oneself; it could also be described as self-loathing. We assume that the short-term self only feels this particular type of disappointment about things within his control; stochastic realizations do not affect his disappointment. So if an agent worked very hard but for some external reason this work was
ineffective, he could rationalize his failure through the external reason, and not feel disappointed. Note that the realized state (influenced by external elements) does not enter this function, but only the action, which is fully controlled by the short-term self. Because this model is mostly geared toward situations of pure selfcontrol, without external commitment devices or control, the only thing preventing one from completing plans is oneself; avoiding disappointing oneself is the main thing that drives plan completion in our model.

At this point in research, we can only speculate about the properties of the disappointment function, as we only begin theorizing about its existence with this paper. It is especially unclear what it "looks like" when its arguments are actions instead of, say, consumption plans. We assume that disappointment is triggered by some distance from a stated goal, but measuring distance over an action space can be a difficult concept. To gain some intuitive clarity, let us think about an action space with a single, continuous numerical dimension, such as hours spent studying or calories consumed, and let us restrict our analysis to actions "worse" than the stated plan, such as studying fewer hours than desired or eating more calories than planned, where the agent would actually experience disappointment. The function is still defined when the agent exceeds expectations, but this is not the realm of focus for this analysis.

Even if we assume disappointment operates on a single numerical dimension, we must also make some assumptions about the shape of the function. Some assumptions are obvious: it is strictly monotonic and positively sloped, that is, it takes negative values for disappointing actions, and more negative values for more actions further away from the goal. We can also assume that there is a discrete jump in the disappointment function, from zero to a negative value, when the realized action moves away from the goal. Comparing disappointment in outcomes
to the better-studied phenomenon of loss aversion in consumption plans opens the possibility of additional assumptions. For one, a standard assumption about utility over losses is decreasing marginal disutility of losses, making the function convex in (effectively) quadrant III. That is, each unit of moving away from the goal is less painful than the last. However, this may not align with psychological realities regarding actions oneself is responsible for, as opposed to outside factors, and it may be concave, especially over smaller absolute values of disappointment. We analyze the effects of different disappointment function in section 5.1.

## Expectations

When determining current expected utility, each agent must take expectations of future possible utilities to make the best current decision. $E_{t, k}$ designates the expectations of the short-term self, given information observed before generating a strategy and the belief system at that information set- that is, $\omega_{t-1}^{*}$ and $\beta_{t, k}$. Similarly, $E_{t, l}$ is the expectation of the long-term self, based on $\omega_{t-1}$ and beliefs $\beta_{t, l}$. The two uncertain elements in the model are the player type, unknown to the opposite agent, and the stochastic element $\epsilon_{\tau}$, unknown until it is realized.

For the long-term player, expectations are formed in the following way: given the long-term self's strategy, a strategy for the short-term self, and a set of beliefs for the long-term player, starting at a specific information set, he can predict the probability of ending up at any particular information set next period, given his plan for that period, the short-term self's plan for that period, his beliefs about the short-term player's type, and the distribution for $\epsilon_{t}$. Then, for all of those resulting information sets in the next period, he can predict probabilities for information sets in the period after that, given strategies and beliefs in those future information
sets. This process continues until the final period is reached, and probabilities are determined for all possible information sets in the final period.

Associated with each of these possible final information sets is a specific discounted utility, from the plans, actions, and state levels that brought the players to that information set. The long-term self can then take the probabilities of all these final information sets, multiplied by their utilities, and form an expected discounted utility, from a current information set, set of strategies for each player, and set of beliefs. The short-term self forms expectations and expected utility in a nearly identical way, but with slightly different information sets.

## Potential Equilibrium Equations

Presentation of the setup of the model is now complete - we have shown and explained information, types, strategies, beliefs, and payoffs for both agents. In this paper, we make no attempt to analyze the existence of an equilibrium in a general setting - the assumptions we have made are too broad for this to be tractable. But if an equilibrium of the model existed, a Weak Perfect Bayesian equilibrium would be the condition of choice: an equilibrium includes optimal strategies for all player types given their beliefs, and updated beliefs over these player types at all information sets.

In an equilibrium, the following best-response conditions would hold. We present them mostly for illustrative purposes, regarding the decision-making processes of the two agents.

$$
\begin{align*}
& \sigma_{k}\left(\omega_{t-1}, a_{t}^{*}, t \mid \theta_{k}, \beta_{k}\right)= \\
& \quad \operatorname{argmax}_{\sigma_{k}} E_{t, k}\left[\sum_{\tau=t}^{T} \gamma_{s, \tau-t}\left(\theta_{k}\right)\left[m_{k}\left(\sigma_{\tau, k}(\bullet)\right)+n_{k}\left(s_{\tau}(\bullet)\right)+D\left(\sigma_{\tau, k}(\bullet) \mid \sigma_{\tau, l}(\bullet)\right)\right] \mid \beta_{k}, \sigma_{l}, \sigma_{k}\right]  \tag{2.11}\\
& \sigma_{l}\left(\omega_{t-1}, t \mid \theta_{l}, \beta_{l}\right)= \\
& \operatorname{argmax}_{\sigma_{l}} E_{l, t}\left[\sum_{\tau=t}^{T} \gamma_{l, \tau-t}\left(\theta_{l}\right)\left[m_{l}\left(\sigma_{\tau, k}(\bullet)\right)+n_{l}\left(s_{\tau}(\bullet)\right)+C\left(\sigma_{\tau, l}(\bullet) \mid \sigma_{\tau, k}(\bullet)\right)\right] \mid \beta_{l}, \sigma_{l}, \sigma_{k}\right] \tag{2.12}
\end{align*}
$$

These equations state that the strategy of the short-term self from a particular point in time onward is a best response to the strategy of the longterm self given a set of beliefs, while the strategy of the long-term self over that same time period is a best responses to the strategy of the short-term self given a set of beliefs. These are fairly obvious equilibrium, but provide some clarity for the thinking behind each agent's decision. Also, in equilibrium, the beliefs of each player must be justified by the original distribution and correct Bayesian updating.

In summary, the model in equilibrium works in the following way: the longterm self picks a plan of action in order to maximize his utility, knowing that the short-term self will react with a realized action in order to maximize his own utility. While he knows the potential strategies of the short-term self given different possible types, because he does not know the real type, he can only predict a distribution of actions in that time period. The short-term self, on the other hand, reacts to that plan of action and picks an action to maximize his utility, discounted
at a different rate than the long-term self, although again taking into account possible future strategies given different types. Both selves trace the effects of their actions into the future into different distributions, then take expectations over all these distributions over time.

### 2.4 Possible Outcomes of the Model

Now that all the elements of the model have been defined and presented, we can discuss how these various elements interact together and the possible results of this interaction. This model has been built in a very general way: we have not assigned any magnitudes to any of our functions, and made few comparisons. The actions predicted by the model depend greatly on the specifics of the model, including shapes of the different utility functions, time preferences, and information.

The simplest case to analyze is where both agents have complete information (that is, $F\left(\Theta_{k} \times \Theta_{l}\right)$ is degenerate), and are able to precisely predict the actions taken by the other agent in response to their own choices. Therefore, the longterm self will only set plans knowing the exact response to those specific plans, and can therefore do a better job of controlling the actions of the short-term self. However, because in our current formulation of the model, the long-term self does not take disappointment into account when maximizing utility, he can possibly choose unrealistic goals in the hope of inducing behavior to avoid even greater disappointment, depending on his consistency cost function. Note that this strategy only works when the slope of the disappointment function is greater than the slope of the disutility of effort.

A more plausible information structure is where each agent does not know the type of the other agent. The most likely application of this is when the long-
term self does not know the degree of present bias that the short-term self has, and is likely to underestimate it in expectation. If the disappointment function is high enough in magnitude, this won't affect the outcome, as the short-term self's benefits to procrastination are outweighed by the disappointment function, making following the plan the resulting strategy. However, if procrastinating is the best strategy for the short-term self, even taking into account disappointment, then unmet plans are possible. We will investigate this possible situation with a numerical example in section 5 .

We briefly mentioned the possibility of extending types to other elements of the utility function; this would make an interesting extension. If we allow the agents to not know the basic utility functions of the other agent, many more complicated results would follow, and the possibilities are too great to analyze here. But one likely case is that the long-term self does not know the shape of the disappointment function. In other words, he does not know how powerful his goals will be in the decision-making process of the short-term self. In this case, even if the long-term self knows the time preferences of the short-term self and wants to set goals to counteract these time preferences, the present-bias of the shortterm self may again be too strong to be counteracted by a weak disappointment function.

### 2.5 Conclusion and Future Work

Combining a number of different theoretical elements, we have created a model with a broad range of possible assumptions to model goal-setting and effort provision in a multi-period setting. In different cases of information awareness between the two agents, we find that behavior found in the real world (ineffective
goal-setting, problematic procrastination, and more) can be produced by the interactions within the model.

This model leaves open the possibility of extensive future work. For one, continuing to refine the model on a purely theoretical level is likely to be productive; in its current state, the model is almost too cumbersome and possibly introduces too many new concepts at once (a disappointment functions, utilities over both actions and states, two sets of full time preferences, etc.). Breaking down the model into smaller components and studying the effects of adding each component to more standard models would be useful and informative.

Generally, the utility and costs of effort and leisure has not been nearly as extensively studied as that of consumption and money, and this model hopes to contribute toward investigating how effort and leisure affect utility in general. Studying the shape of the different utility components (especially including the cost of effort and the disappointment function) with experimental techniques could be rewarding. Also, experimentally studying time preferences as related to effort in both a long-term state of mind and a short-term state of mind could be very revealing. While some of this has been done before, continuing to study self-awareness of tendency to procrastinate, especially in a more controlled setting, would be especially helpful in working toward a realistic set of assumptions for this model.

We hope that the structure and assumptions set up in this model can help economists better understand human behavior of goal-setting and effort provision, and how real-life behavior defies traditional logic, but may do so in unexpectedly structured ways.

## CHAPTER III

## ESTIMATING TIME PREFERENCES FOR SCHOOL WORK USING CONVEX TIME BUDGETS

Time preferences for money have been extensively studied, but experimental studies of time preferences for effort and leisure are rare. Using a field experiment embedded in an introductory economics class, we studied time preferences for completing multiple-choice review questions for graded credit. Through an online survey students were given convex time budgets with varying time durations, intertemporal exchange rates, and numbers of questions, and selected how many review questions to complete during given dates, and then completed the questions online. Reduced form analyses and aggregate structural results show that many student subjects have strong preferences for allocations that evenly divide work across time periods. Many others are effectively debt-averse - they choose to complete questions sooner, even when that results in more total work. Some expected outside effort and time commitments (class attendance, jobs, social events, etc.) had small and marginally significant effects on choices. Incomplete task completion may indicate subjects were attempting to implement commitments and failing, and that the experiment may not have correctly elicited true time preferences.

### 3.1 Introduction

Time preferences are one of the most well-studied topics in behavioral economics, with many different experimental techniques developed to analyze discount rates, functional forms of the discount function, degrees of present-bias,
and more. However, most previous work has only investigated time preferences for money. Time preferences for things other than money, like consumption or leisure, could be very different. The literature on procrastination, which often is in reference to effort versus leisure, rarely relates to experimental literature, nor does it acknowledge that individuals could have different time preferences in different domains. Also, very little has been done to investigate how much outside factors affect specific decisions about allocations over time, such as outside income or expenditures, or outside effort provision.

In this experiment, we use a recently-developed experimental technique (Convex Time Budgets) to investigate time preferences that students have for completing school work. As part of an Introduction to Microeconomics class, this experiment had students allocate the task of completing multiple-choice review questions between different time periods with different exchange rates between questions earlier and questions later, to investigate how students chose to allocate effort under different circumstances.

Results show that in aggregate, students responded to differences in prices, but not significantly to differences in how far in the future the tasks would occur. However, there was considerable variation in how individuals behaved, partially caused by outside time commitments and variation in the date on which the experiment was completed, but also from apparent variation in time preferences.

Section 2 reviews the relevant literature on the topic, including studies about time preferences, real-effort tasks, and procrastination. Section 3 describes the experimental procedure, while section 4 presents the empirical strategies employed in analysis. Section 5 shows results from the population, including qualitative and quantitative results, and Section 6 concludes.

### 3.2 Literature Review

The study of time preferences can be divided along two different criteria into four categories of preferences: gains versus losses, and monetary rewards versus primary consumption rewards. (Note that while economists use the phrase "time preferences", psychologists tend to use the phrase "delay discounting" to refer to effectively the same thing.) This experiment will focus on time preferences for losses of primary consumption - specifically, the loss of leisure through real-effort tasks. This particular combination has received very little focus in the literature, yet other areas related to it (time preferences for monetary gains and losses, and consumption gains) have received large amount of attention, particularly in recent years (Hardisty et al. (2013) contains a graph showing a dramatic increase in research volume on the subject).

There have been proposed many mathematical models of time preferences, including most prominently hyperbolic discounting (first proposed by Kahneman and Tversky (1979)) and quasi-hyperbolic discounting (from Laibson (1997)). Doyle (2013) surveys all the various models of time preferences, and the underlying theories and assumptions that generate those models. This paper will align with much of the rest of the current economic literature, and use Laibson's quasi-hyperbolic or beta-delta preferences, due to its previous empirical success, theoretical simplicity and ease of use.

The vast majority of the economics literature on time discounting has examined discounting of monetary gains - see Frederick et al. (2002) for a review of earlier work. Because of the focus this area receives, experimental and econometric techniques have developed most rapidly here. More recently, Andersen et al. (2008) investigate risk and time preferences in a representative sample of the

Danish population through a Double Multiple Price List (DMPL) procedure, while Andreoni and Sprenger (2012a) use Convex Time Budgets (CTB) in an attempt to estimate utility curvature separately from risk preferences. These two approaches, DMPLs and CTBs, are currently the most favored techniques for eliciting time preferences from individuals while controlling for utility curvature. Both sides have criticized the approaches of the other: the DMPL researchers have stated that the results from CTBs do not align with how individuals should act according to their assumptions (see Harrison et al.), while the CTB researchers have stated that using risk preferences to elicit utility curvature related to time preferences attempts to equate two unrelated sets of preferences (see Andreoni and Sprenger (2012b)). See Andreoni et al. (2013) for a theoretical and econometric comparison of the two approaches. Hardisty et al. (2013) find different methods with the same subjects result in slightly different results.

The research on discounting of monetary rewards generally acknowledges that individuals receive utility not from money but from the consumption that money brings, and therefore studies of money discounting only provide an approximation for time preferences of utility. Also, external capital markets exist for money, allowing for arbitrage to influence experimental findings. Cubitt and Read (2007) discuss this connection and present different theories behind the connection between monetary tasks and utility discounting and a theoretical model for dealing with these differences. While monetary tasks can shed some light on discounting of utility, they are certainly less than perfect.

Given this criticism, a few researchers have investigated time preferences for direct consumption through giving primary rewards - consumables that are more closely related to consumption utility at a particular time than money. The famous Stanford Marshmallow Experiment (discussed in Mischel et al. (1989)) began to
investigate the delay of consuming goods (or primary rewards), which essentially reveals time preferences. McClure et al. (2007) investigate the neuroeconomics of primary rewards, finding the neurological pathways related to primary rewards are similar to those of monetary rewards, but that even a ten minute time delay between decision and reward (in this case sips of juice or water) can reduce stimulation of the limbic system (the brain's short-term pleasure system) when making decisions. However, the nature of these small rewards makes scaling over time and volume difficult. Estle et al. (2007) compare discount rates for directly consumable goods, including candy, soda and beer, and finds that these goods are discounted more steeply than monetary rewards, but that this difference disappears when results are probabilistic. The utility from monetary gains is assumed to be a conditioned response to anticipation of future rewards, rather than the direct consumption utility of primary rewards. Reuben et al. (2010) find similar results using real (rather than hypothetical) results.

However, these results are confounded by issues of magnitude and diminishing marginal utility. Another problem is that demand for primary rewards can be greatly affected by the state of mind of the individual regarding the rewards. We know that time preferences for monetary rewards depends on the state of mind of the individual, especially through exitation of pleasure-seeking neural pathways, (see Van den Bergh et al. (2008) for an entertaining example), but appetite for the particular primary reward is much more variable than appetite for more money (as Reuben et al. (2010) discuss).

Along with differences between monetary versus primary rewards, time preferences can also differ depending on if they are over gains or losses. Prospect theory (from Kahneman and Tversky (1979)) assumes losses produce disutility larger in magnitude to the utility from corresponding gains, but time preferences
for losses can be completely different. Monetary losses exhibit some similar elements of discount compared to gains - for example, Holt et al. (2008) find the existence of preference reversals for losses, which implies a hyperbolic or quasihyperbolic discount function, resulting in present-bias. However, there are many differences in discounting of gains versus losses - Estle et al. (2006) show many of the differences, including a smaller or even reversed magnitude effect for losses compared to gains, but closer results for probabilistic gains and losses. Appelt et al. (2011) create psychological explanations for these differences in discounting behavior between gains and losses. Frederick and Loewenstein (2008) look not specifically at time preferences, but at sequences of events, and investigate whether subjects have preferences for worsening, improving, or flat sequences of events of real-world outcomes, and find varied results.

Prelec and Loewenstein (1998) also shows that mental accounting of debt and payments is different from accounting of savings, and that individuals often prefer to pay off debt or pay for consumption as soon as possible, even if it runs counter to their economic benefit. However, much of the research on losses has been with purely hypothetical rewards - in experiments, real losses (that is, forcing subjects to give up their own money) must be compensated either beforehand through endowments or afterward through rewards, both of which can complicate behavior. In forthcoming work, Kuhn, Andreoni and others investigate discounting behavior of real losses and find some counterintuitive results, suggesting that subjects view debt very differently, possibly suggesting concave preferences for losses.

Investigating losses of primary rewards, however, is even more difficult. Taking away consumption goods from experimental subjects is effectively impossible unless they are endowed with them in the first place, which creates problematic interactions with reference points. However, forcing subjects to expend
leisure time on a task is one way of having subjects give up a consumption good (leisure). Also, if the tasks involved include cognitive effort, this could potentially cause disutility - Westbrook et al. (2013) investigate preferences for cognitive effort and find subjects are willing to forgo monetary gains to avoid increased cognitive effort tasks, although in other circumstances cognitive tasks may be enjoyable. Therefore, we can view decreased leisure time and increased cognitive load as effectively a loss of consumption utility. The purpose of our experiment is to investigate time preferences for this disutility of effort generated by giving up leisure and increasing cognitive load in a task that most rational people agree is unpleasant - answering questions about economics!

Our experiment most closely resembles the experimental work of Augenblick et al. (2015), who examine time preferences and commitment devices regarding real-effort tasks. Subjects in this study allocate real-effort tasks, specifically playing a modified Tetris game and transcribing fuzzy Greek letters, between two time periods given different exchange rates. This is one of very few examples of estimating time preferences for losses of primary rewards. However, the authors do not present their findings as losses compared to gains of consumption utility, and do not discuss the different aspects that preferences over losses can take versus preferences for gains. The tasks included in the experiment were also basic, requiring minimal cognitive effort, so subjects were only choosing to give up leisure time instead of having increased cognitive effort and its associated disutility. Also, while the experiment was designed to find present-bias and associated commitment costs that could possibly reverse present-bias, the design only uses time scales of one week between decisions, making it difficult to estimate discounting. Our experimental procedure will be designed to more accurately measure discount rates over slightly longer time periods and the degree of present-bias involved in losses
of consumption. Augenblick and Rabin (aper) also investigate time preferences for real-effort tasks, but use different monetary wages for tasks to investigate time preferences. They also find significant evidence of present-bias, but their results also don't include outside effort, nor address the confounding effects of money.

Another paper using a similar real-effort task is Bisin and Hyndman (2014b) - subjects must alphabetize nonsense words as a real-effort task. However, the purpose of this experiment is not to specifically measure time preferences, but rather find evidence of present-bias and how it effects procrastination behavior. While procrastination behavior is the most obvious effect of present-biased time preferences for real-effort tasks, measuring the degree of present-bias and discounting that create this procrastination behavior is an important task.

### 3.3 Experimental Procedure

Our experimental procedure borrows heavily from Augenblick, Niederle and Sprenger (2014), in using convex time budgets to estimate time preferences for realeffort tasks, but introduces a number of new elements, including controlling for opportunity costs, applying it to real-world academic applications, and attempting to correlate our results with other indications of present-biased behavior.

Participants were recruited from a large Principles of Microeconomics course through an in-class introduction to the experiment. Students could choose to either participate in the experiment, including answering the included review questions, or write a 5-page essay on microeconomics based on news articles. This alternative assignment and its deadline structure may have created some selection bias, but considering that the alternative assignment was more difficult and took more time, and over $85 \%$ percent of students did participate, selection is not a major issue.

After the in-class introduction and after taking the midterm exam, students completed the experimental choices in an online survey, taken on a day of the week of the students choice. As described in the introductory session and other recruitment materials, subjects could choose which day of the week to complete their preference choices and subsequent real-effort tasks. Students were also informed that they would be completing tasks on the day they made their allocations, two weeks from that date, and four weeks from that date. This allowed subjects to plan around their outside commitments, making the experiment as little of a hassle as possible, increasing participation, and eliciting responses based on time preferences, not responses based on outside commitments. This allocation process took place in the week after the midterm exam in the class.

The real-effort tasks in this experiment were multiple-choice review questions drawn from the test bank of the microeconomics textbook used by that class. These review questions were all based on subject material covered on the midterm exam, and that would not be explicitly covered by the final exam (which only covered material from the second half of the course). Therefore, there was no incentive to complete questions closer to the final exam to be more prepared for it, but there was an incentive to complete them closer to the midterm, as the material was fresh in their minds. At the beginning of the online survey, subjects were shown two examples of questions to help them gauge their difficulty and the amount of time necessary to complete them. Then, subjects entered the number of hours of outside commitments for each day they could complete questions, which at the time were that current day, two weeks from that current day and four weeks from that day (described as calendar dates, not relative times) from the subjects perspective. These outside commitments included classes, study time, social functions, and other categories. See figure 1 for an example (with two
categories cut off). Subjects also rated their responses about commitments by level of confidence.

In the primary section of the experiment, subjects allocated review questions between two dates using Convex Time Budgets (CTBs), from Andreoni and Sprenger (2012a). Subjects made nine different allocations, one for each date combination (three combinations of that present day, in two weeks, and in four weeks), and one for each review question exchange rate (3:5, 5:5, and 5:3). For example, given the 3:5 exchange rate, subjects were endowed with 35 questions for the earlier time period, and could exchange 3 questions in the earlier period for 5 questions in the later, with eleven possible options. These options were presented in a multiple-choice list. Subjects faced a minimum of five questions for each time period within each CTB, to equalize transaction costs of completing tasks across time periods. Due to the quantized nature of review questions versus money and possible confusion about the options, this was chosen as the best presentation option, as subjects also were able to view all of their different resulting options immediately without doing any calculations. See figure 2 for an example.

To compare our results to previous work, students also completed CTBs for allocating money between the same dates as they would complete review questions. However, these allocations were only hypothetical, for budgetary and logistic reasons, and possibly biased by the possibility of having to complete tasks on the given days. Results from these decisions are not reported, as they show minimal subject investment in those decisions.

For their final questions, subjects answered a number of basic demographic questions, including year and major, and gave a self-rating about behaviors related to procrastination, both in general (taken from Lay (1986)) and specific to
Think about your time commitments in week 6.


FIGURE 1. Commitments (missing Family and Social time)

Here, allocate questions between your chosen date in week 8 , and your chosen date in week 10 . The exchange rate for questions is $5: 3 ; 5$ questions in week 8 are worth 3 questions in week 10.

$$
\bigcirc
$$ 55 questions on Wednesday, April 22nd and 5 questions on Wednesday, May 6th

50 questions on Wednesday, April 22nd and 8 questions on Wednesday, May 6th

- 45 questions on Wednesday, April 22nd and 11 questions on Wednesday, May 6th
$\qquad$ 40 questions on Wednesday, April 22nd and 14 questions on Wednesday, May 6th

35 questions on Wednesday, April 22nd and 17 questions on Wednesday, May 6th 30 questions on Wednesday, April 22nd and 20 questions on Wednesday, May 6th 25 questions on Wednesday, April 22nd and 23 questions on Wednesday, May 6th 20 questions on Wednesday, April 22nd and 26 questions on Wednesday, May 6th - 15 questions on Wednesday, April 22nd and 29 questions on Wednesday, May 6th 10 questions on Wednesday, April 22nd and 32 questions on Wednesday, May 6th - 5 questions on Wednesday, April 22nd and 35 questions on Wednesday, May 6th

FIGURE 2. Allocation Decision
academic settings. These were used to investigate perceived real-world consequences of having specific time preferences.

After entering their allocations and completing the survey questions, one of the allocations of review questions was chosen at random for implementation. Because each allocation covered two time periods, subjects were assigned questions on two out of the three possible time periods. For questions to be completed that day, subjects continued within the same online survey to the review questions to be completed; for questions to be answered in the future, subjects received an email link to the assigned questions early in the morning on the day on which they had to be completed. Once all rounds of questions were completed, the percentage of assigned questions answered correctly was submitted to the instructor of the course to count for $5 \%$ of the final grade in the class.

### 3.4 Empirical Strategy

## Reduced Form Analysis

In the first stage of our analysis, we analyze a reduced-form model, testing how subjects respond to differences in time, price, and outside effort, and how different responses to demographic and survey questions influence choices of effort. In this analysis, we use the number of questions allocated to the earlier period as the dependent variable. There are a number of different empirical approaches to this estimation, though.

Because subjects cannot choose fewer than five questions for any particular time period, the data is effectively truncated, both above and below. The coefficients in the baseline OLS model are therefore biased toward zero. The traditional way of dealing with data censoring, the Tobit model, can resolve some
of this truncation problem. However, due to the experimental design, some data is truncated above at 55 , while other data is truncated at 35 . The Tobit model can only deal with one truncation point in each direction, so it would also return biased results.

The solution to this issue is the Interval Regression, which is effectively a generalized form of the Tobit model. This method uses two dependent variables, an upper and lower bound for each data point, and estimates regressions using the following likelihood function, where $U^{*}$ is the relevant upper limit.

$$
\begin{equation*}
L=\prod_{x_{t}=L}\left(1-\Phi\left(\frac{x_{t}^{*}-L}{\sigma}\right)\right) * \prod_{L<x_{t}<U^{*}}\left(\frac{1}{\sigma} \phi\left(\frac{x_{t}-x_{t}^{*}}{\sigma}\right)\right) * \prod_{x_{t}=U^{*}}\left(\Phi\left(\frac{x_{t}^{*}-U^{*}}{\sigma}\right)\right) \tag{3.1}
\end{equation*}
$$

Because the data is quantized by 5 or 3 questions, depending on the exchange rate, we can also use an interval regression and assume that each data point only indicates an interval between midpoints of choices, instead of a point observation at each choice. The likelihood function for that is the following, where $x_{1 t}$ and $x_{2 t}$ are the upper and lower bounds for each observation:

$$
\begin{equation*}
L=\prod_{x_{t}=L}\left(1-\Phi\left(\frac{x_{t}^{*}-L}{\sigma}\right)\right) * \prod_{L<x_{t}<U^{*}}\left(\Phi\left(\frac{x_{2 t}-x_{t}^{*}}{\sigma}\right)-\Phi\left(\frac{x_{1 t}-x_{t}^{*}}{\sigma}\right)\right) * \prod_{x_{t}=U^{*}}\left(\Phi\left(\frac{x_{t}^{*}-U^{*}}{\sigma}\right)\right) \tag{3.2}
\end{equation*}
$$

## Structural Analysis

Following most previous work in time preferences, we assume that agents intertemporally optimize their utility, based on a utility function. We assume that this utility function is time-separable, stationary, and satisfies the independence
axiom (so that we can ignore the probabilities of different options being selected). In our case, we are looking at (dis-)utility of effort, which we assume is a combined measure the mental or physical energy put into a task, and the time spent doing so. For the functional form of the utility function, here effort in one time period is $e_{t}$, we assume it is:

$$
\begin{equation*}
u\left(e_{t}\right)=-\frac{1}{\alpha} e_{t}^{\alpha} \tag{3.3}
\end{equation*}
$$

Here, $\alpha$ is the curvature of disutility of effort. $\alpha>1$ implies increasing marginal cost of effort (or decreasing marginal utility of leisure), while $\alpha<1$ implies decreasing marginal cost of effort. We assume that agents maximize their time-discounted utility (or minimize time-discounted costs) given the budget restriction of having to complete a certain number of tasks. Unlike previous studies, we assume that effort includes both effort related to the task, and effort related to outside activities, such as studying, attending class, or participating in other outside activities. So $e_{t}=x_{t}+\omega_{t}$, where $x_{\tau}$ is the number of questions allocated in the experiment, and $\omega_{\tau}$ is the total outside effort predicted by the subject. Therefore, agents have the following problem:

$$
\begin{equation*}
\max U_{t}=\sum_{\tau=0}^{n} \gamma_{\tau}-\frac{1}{\alpha}\left(x_{\tau}+\omega_{\tau}\right)^{\alpha} \tag{3.4}
\end{equation*}
$$

Importantly, we do not know how much effort specific outside tasks require, compared to the effort involved in completing the review questions. Therefore, we assume each outside task $\eta$ has a weight $\theta$, relative to review questions:

$$
\begin{equation*}
\omega_{t}=\lambda+\sum_{i=1}^{n} \theta_{i} \eta_{i} \tag{3.5}
\end{equation*}
$$

We also assume that subjects have quasi-hyperbolic or beta-delta time preferences:

$$
\gamma_{t}= \begin{cases}1 & t=0  \tag{3.6}\\ \beta \delta^{t} & t>0\end{cases}
$$

Note that unlike the traditional beta-delta model with preferences over gains, our preferences are over negative utility. We know that individuals are often debtaverse, and that promises of future utility losses are especially painful. Therefore, it may be possible that $\beta>1$, which would signify subjects being debt-averse.

The Convex Time Budget experimental procedure gives subjects a budget under which they must maximize their utility. So, under this budget and given all of our utility assumptions, subjects have to solve the following problem in each task, over two periods:

$$
\begin{align*}
& \max _{x_{\tau}}-\beta^{t \neq 0} \delta^{t}\left(x_{t}+\sum \theta_{i} \eta_{i, t}\right)^{\alpha}-\beta \delta^{t+k}\left(x_{t+k}+\sum \theta_{i} \eta_{i, t+k}\right)^{\alpha}  \tag{3.7}\\
& \quad \text { s.t. } P_{1} x_{t}+P_{2} x_{t+k}=m \tag{3.8}
\end{align*}
$$

## Nonlinear Least Squares

Our first estimation approach utilizes Nonlinear Least Squares (NLS) to estimate the parameters of the model. For this approach, we analytically maximize utility and find the following condition:

$$
\begin{equation*}
x_{t}=\frac{\left(\frac{m}{P_{2}}+\omega_{t+k}\right)\left(\frac{\gamma_{t+k} P_{1}}{\gamma_{t} P_{2}}\right)^{\frac{1}{(\alpha-1)}}}{\left(1+\frac{P_{1}}{P_{2}}\left(\frac{\gamma_{t+k} P_{1}}{\gamma_{t} P_{2}}\right)^{\frac{1}{(\alpha-1)}}\right)}-\frac{\omega_{t}}{\left(1+\frac{P_{1}}{P_{2}}\left(\frac{\gamma_{t+k} P_{1}}{\gamma_{t} P_{2}}\right)^{\frac{1}{(\alpha-1)}}\right)} \tag{3.9}
\end{equation*}
$$

Substituting in the parameters to be estimated as defined above $\left(\beta, \delta, \alpha, \theta_{i}, \lambda_{i}\right)$, and the data ( $m, P_{1}, P_{2}, \eta_{i, \tau}$ ), and assuming an additive, normally distributed error, we can estimate the above equation using the standard NLS procedure.

This estimation strategy is flexible in its utility curvature assumptions: if $\alpha>1$, subjects have increasing marginal disutility of effort (each task is worse than the previous one), and will seek interior solutions on a budget. If $\alpha<1$, there is decreasing marginal disutility (additional tasks become more tolerable), and subjects will seek corner solutions.

However, a weakness of the standard NLS proceduce is its inability to deal with the truncating issues related to corner solutions, which produces biased estimates. Simulations show this bias is especially strong for the $\theta_{i}$ parameters. Because subjects cannot assign fewer than five questions to either day in each allocation, but may want to assign fewer than that or even negative questions, we will observe corner solutions on both the upper and lower bounds of the allocation of $x_{t}$. Future work will implement methods to mitigate the bias that truncation creates.

It may be that subjects view effort in different time periods differently, as questions in different periods may require different amounts of mental effort, or the expected grade payoff is different for them, accounting for the possibility of
forgetting to complete questions. Therefore, we include two possible variations on the estimating equation. In the first, baseline cost of effort changes additively across different time periods: $\lambda_{t} \neq \lambda_{t+1}$. Second, we allow multiplicative scaling of effort between different periods: effort $=\left(x_{t}+\omega_{t}\right)\left(1+z_{t}\right)$, with the normalizing assumptions that $z_{0}=0$. This allows subjects to assume different payoffs for questions in different periods, or different cost of effort in different periods. While varying $\alpha$ between periods (allowing the curvature of utility to change) would be another possible route, allowing this makes the maximization problem algebraically impossible in a closed-form solution.

## Alternative Perspective

The previously described assumptions about utility are that subjects derive disutility from effort, compared to a zero-effort baseline. However, another possible perspective about subject's utility is that they gain utility from leisure, and effort represents a decrease leisure. Note that we are not discussing "gains" versus "losses" of leisure compared to a reference point in the Prospect Theory sense, as our experimental procedure does not compare gains versus losses. We are just assuming that in this frame, agents maximize positive utility of leisure, instead of minimizing disutility of effort. If we assume that $\lambda$ is a daily leisure budget instead of baseline daily effort, and effort subtracts from this daily leisure budget to form total leisure $l_{t}$, the utility formulation would be:

$$
\begin{equation*}
u\left(e_{t}\right)=\left(l_{t}\right)^{\alpha^{\prime}}=\left(\lambda-e_{t}\right)^{\alpha^{\prime}} \tag{3.10}
\end{equation*}
$$

The consequences of different utility curvature are very different given this perspective. Here, $\alpha^{\prime}<1$ implies decreasing marginal utility of leisure (a standard
assumption about goods in general, and within labor economics), which leads to interior solutions, while $\alpha^{\prime}>1$ generates increasing marginal utility of leisure (counter to many assumptions, but not impossible) and creates corner solutions. (Compare to the curvature of the effort disutility function: $\alpha>1$ implies interior solutions, and vice versa.)

The analytic maximization results are very similar, although now making explicit the $\lambda$ term which was previously included in the $\omega_{t}$ term:

$$
\begin{equation*}
x_{t}=\frac{\left(\frac{m}{P_{2}}+\omega_{t+k}-\lambda\right)\left(\frac{\gamma_{t+k} P_{1}}{\gamma_{t} P_{2}}\right)^{\frac{1}{(\alpha-1)}}}{\left(1+\frac{P_{1}}{P_{2}}\left(\frac{\gamma_{t+k} P_{1}}{\gamma_{t} P_{2}}\right)^{\frac{1}{(\alpha-1)}}\right)}-\frac{\omega_{t}-\lambda}{\left(1+\frac{P_{1}}{P_{2}}\left(\frac{\gamma_{t+k} P_{1}}{\gamma_{t} P_{2}}\right)^{\frac{1}{(\alpha-1)}}\right)} \tag{3.11}
\end{equation*}
$$

We can then use our previous methods of truncated NLS or Tobit estimation to calculate all the same parameters, but including $\lambda$. This formulation can be better compared with results for discounting of positive monetary income, as it assumes a closer utility functional form to those previous studies.

### 3.5 Descriptive and Aggregate Results

## Descriptive Results

Figure 3 shows the means of these allocation decisions shown on the budget sets restricting the decision. We can see that the means for each time period combination are nearly identical, and that the mean allocation is very close to the middle of the given budget. Figure 4 shows a histogram of the number of review questions selected for the first period in each of the nine decisions; Price

FIGURE 3. Allocation within Budgets

is the exchange rate between questions in the first period versus the second price $=.6$ means three questions in the first period are worth five in the second, while price $=1.667$ means five questions in the first period are worth three in the second. Note that in the left column of graphs, the maximum number of questions available in the first period is 35 , while in other columns it is 55 . From this graph, we can draw a number of conclusions about the distribution of preferences from visual inspection.


Weeks $8 \& 10$, price $=1.667$
mean $=29.48$





FIGURE 4. Review Question Allocation

One obvious trend in this data is the similarity of choice distributions across time periods - there is very little apparent change when comparing decisions between weeks $6 \& 8$ versus weeks $6 \& 10$ or weeks $8 \& 10$. This suggests that time preferences play a smaller role in subject decision-making than prices and rules-ofthumb, or that time preferences are not very strong.

Another prominent feature is that given an even exchange rate, individuals have a strong preference for dividing work evenly - over half of subjects consistantly choose an even split given an even exchange rate. This suggests minimal discounting, but also possibly a reliance on rules-of-thumb when allocating tasks.

A main topic of research in this experiment is the curvature of the value funciton for real-effort tasks - $\alpha$ in our model. As discussed, $\alpha>1$ implies increasing marginal cost of effort, while $\alpha<1$ implies decreasing marginal cost of effort. With $\alpha>1$, subjects would likely have interior solutions for the maximization problem - they seek to balance their effort between the two time periods until the marginal cost of effort is equal. But when $\alpha<1$, they would likely have corner solutions - because each subsequent question gets easier, they would want to complete as many as possible at one time. Within our results, we find a number of corner solutions and a number of interior solutions, suggesting a diversity of values for $\alpha$ for individuals.

Thinking about time preferences within each allocation decision, we find that outside of the even exchange rate, many subjects seem to have a preference for completing tasks sooner rather than later. Even when they could complete fewer total questions by putting them off, many subjects choose to complete more questions in order to do them sooner. This is especially apparent in the right column, where doing more questions now results in more total questions. This suggests that some subjects "preproperate." This word, meaning to do too
soon and serving as the opposite of procrastinate, is rarely used, but describes the situation for some subjects well.

There are a number of possible behavioral explanations for this preproperation. One is that subjects have a preference for improving sequences of events - they want to get unpleasant tasks over with quickly, and therefore are happier in the future. Subjects also might prefer to complete the review questions while the class material on which they were based is fresh in their minds, or they may fear forgetting to complete review questions on later dates.

## Reduced Form Results

Just as in the histograms, our dependent variable of interest is the number of questions allocated to the first period. Table 1 shows the reduced form estimation for the various different structural parameters - Interval assumes most data is point data, other than end points, while Interval2 assumes all data is interval data between midpoints of selection points. Present indications whether or not the comparison includes present bias (that is, includes the current period), Weeks is the number of weeks between time periods for each allocation, Exchange is the exchange rate between questions in the first period versus second, and Endowment is the size of the initial allocation (which depends on the exchange rate). ClassA, etc., are the time commitments in each category in the first and second date of each allocation. Note that negative coefficients mean fewer questions completed earlier, indicated increased procrastination, or decreased "preproperation". All of our reduced-form equations include clustering of standard errors at the individual level.

We can see that in comparing between allocations, the coefficients Present and Weeks are insignificant, and in fact are fairly tight zeros. This means that

TABLE 1. Reduced Form

|  | $(1)$ | $(2)$ | $(3)$ | $(4)$ |
| :--- | :---: | :---: | :---: | :---: |
|  | OLS | Tobit | Interval | Interval2 |
| Exchange | $-7.851^{* * *}$ | $-8.814^{* * *}$ | $-8.892^{* * *}$ | $-8.588^{* * *}$ |
| Endowment | $(-8.75)$ | $(-8.27)$ | $(-8.26)$ | $(-8.47)$ |
|  | $0.520^{* * *}$ | $0.572^{* * *}$ | $0.523^{* * *}$ | $0.517^{* * *}$ |
| Present | $(19.30)$ | $(16.88)$ | $(15.47)$ | $(16.35)$ |
|  | -0.121 | -0.0406 | 0.0108 | -0.0386 |
| Weeks | $(-0.20)$ | $(-0.06)$ | $(0.02)$ | $(-0.06)$ |
|  | 0.0495 | 0.0461 | 0.0324 | 0.0408 |
|  | $(0.21)$ | $(0.17)$ | $(0.11)$ | $(0.15)$ |
| assignA | $0.600^{*}$ | 0.667 | 0.701 | 0.672 |
|  | $(1.73)$ | $(1.61)$ | $(1.58)$ | $(1.62)$ |
| assignB | 0.323 | 0.381 | 0.350 | 0.338 |
|  | $(1.01)$ | $(1.01)$ | $(0.87)$ | $(0.89)$ |
| workA | -0.589 | -0.785 | -0.798 | -0.741 |
|  | $(-1.40)$ | $(-1.49)$ | $(-1.40)$ | $(-1.39)$ |
| workB | 0.548 | 0.696 | 0.723 | 0.686 |
|  | $(1.35)$ | $(1.32)$ | $(1.28)$ | $(1.30)$ |
| familyA | $-0.600^{* *}$ | $-0.721^{* *}$ | $-0.771^{* *}$ | $-0.726^{* *}$ |
|  | $(-2.42)$ | $(-2.35)$ | $(-2.28)$ | $(-2.30)$ |
| familyB | $0.961^{* * *}$ | $1.170^{* * *}$ | $1.246^{* * *}$ | $1.160^{* * *}$ |
|  | $(2.78)$ | $(2.78)$ | $(2.60)$ | $(2.60)$ |
| socialA | -0.263 | -0.363 | -0.304 | -0.270 |
|  | $(-0.88)$ | $(-1.02)$ | $(-0.79)$ | $(-0.75)$ |
| socialB | $-0.544^{*}$ | $-0.604^{*}$ | $-0.689^{*}$ | $-0.663^{*}$ |
|  | $(-1.84)$ | $(-1.74)$ | $(-1.84)$ | $(-1.88)$ |
| cons | $8.576^{* * *}$ | $6.933^{* * *}$ | $10.11^{* * *}$ | $9.963^{* * *}$ |
|  | $(6.33)$ | $(4.32)$ | $(5.71)$ | $(6.01)$ |
| $N$ | 2684 | 2684 | 2684 | 2684 |
| $t$ statistics in | parentheses |  |  |  |
| $* p<0.10,{ }^{* *}$ | $p<0.05,{ }^{* * *} p<0.01$ |  |  |  |

allocations for comparing periods 1 and 2 versus periods 1 and 3 or periods 2 and 3 are effectively the same, which suggests that in aggregate, subjects do not respond differently to different time periods, at least in our reduced form analysis. There are a number of possible explanations for this: first, subjects are paying much more attention to the exchange rate than to the time periods - we do see that the exchange rate and initial allocation are significant. Second, subjects could only be thinking about "sooner" versus "later", effectively viewing all comparisons as between the same two periods. This attitude would completely hide time preferences from this reduced form analysis. Note, however, that structural analysis, by allowing for curvature of utility, would be able to identify time preference parameters from this data.

We find that most of the coefficients on outside effort are not significant subjects do not seem to take into account their outside time commitments when allocating questions. However, the time commitments data is very noisy - subjects were not incentivized to answer correctly, and many gave unrealistic answers. Also, many subjects were not confident in their commitments, especially for dates a month in the future. In an attempt to reduce bias, we re-run the Interval2 regression with various sample restrictions: Realistic includes only subjects who listed fewer than 24 hours of commitments on each date, Confident only includes subjects in the upper third of rated confidence in responses, and RealConf applies both restrictions. Table 2 shows these results. Given these restrictions, a few of the outside effort coefficients become significant, suggesting measurement error was a problem in the first regression, and that some of these outside commitments do make a small difference in decision-making.

Using our reduced form results, we can also investigate how demographics affect allocations. Table 3 shows the results for a number of demographic

TABLE 2. Sample Restrictions

|  | (1) <br> FullSample | $\overline{(2)}$ <br> Realistic | (3) <br> Confident | (4) <br> RealConf |
| :---: | :---: | :---: | :---: | :---: |
| Exchange | -8.588*** | -9.013*** | -8.850*** | -9.311*** |
|  | (-8.47) | (-8.08) | (-5.76) | (-5.72) |
| Endowment | $0.517^{* * *}$ | $0.535^{* * *}$ | 0.545*** | 0.569*** |
|  | (16.35) | (14.33) | (11.40) | (10.59) |
| Present | -0.0386 | 0.0461 | -0.641 | -0.915 |
|  | (-0.06) | (0.06) | (-0.66) | (-0.90) |
| Weeks | 0.0408 | 0.0408 | 0.854* | 0.790 |
|  | (0.15) | (0.12) | (1.93) | (1.62) |
| assignA | 0.672 | 0.491 | 1.513** | 2.000** |
|  | (1.62) | (0.84) | (2.33) | (2.39) |
| assignB | 0.338 | 0.877* | -0.143 | -0.356 |
|  | (0.89) | (1.75) | (-0.22) | (-0.47) |
| workA | -0.741 | -0.721 | -1.319* | -1.434* |
|  | (-1.39) | (-0.97) | (-1.76) | (-1.82) |
| workB | 0.686 | 0.394 | 1.124 | 1.175 |
|  | (1.30) | (0.51) | (1.63) | (1.50) |
| family A | -0.726** | -0.684 | 0.0722 | 0.191 |
|  | (-2.30) | (-1.53) | (0.14) | (0.28) |
| familyB | $1.160^{* * *}$ | 1.501** | 0.952* | 1.076 |
|  | (2.60) | (2.52) | (1.65) | (1.47) |
| socialA | -0.270 | -0.621 | -0.895* | -0.893 |
|  | (-0.75) | (-1.33) | (-1.67) | (-1.26) |
| socialB | -0.663* | -0.414 | 0.111 | 0.215 |
|  | (-1.88) | (-0.83) | (0.18) | (0.27) |
| _cons | 9.963*** | $7.976{ }^{* * *}$ | 6.349** | 4.572 |
|  | (6.01) | (3.41) | (2.55) | (1.46) |
| $N$ | 2684 | 2205 | 1143 | 999 |
| $t$ statistics in parentheses |  |  |  |  |
| ${ }^{*} p<0.10,{ }^{* *} p<0.05,{ }^{* * *} p<0.01$ |  |  |  |  |

categories; this includes controls for all variables seen in Table 1. An interesting pattern is that as students get older, they seem to put off tasks more - this could be from increased burnout leading to increased procrastination, or increased sophistication and self-trust, leading to less preproperation. There has been significant speculation that Chinese language speakers may exhibit different time preferences from English speakers, as their languages deals with time and tense very differently, but we find no difference between them. Other language differences are driven by small sample sizes.

Table 4 shows whether survey questions about real-life general and specific behaviors are related the number of questions subjects allocate to the first period. These questions were all asked after allocation decisions had been made, so there was no chance of subjects keeping these questions in mind while answering questions. Note that the first seven survey questions are measured on a 5 -point Likert scale, with 5 as the "Strongly agree". The final four behavior questions are binary yes/no questions. Most of the survey questions are insignificant, but the survey question "I often have a task finished sooner than necessary" does indicate that subjects complete more questions sooner (note the asterisk indicates a reversal of the variable, so all coefficients would be expected to be positive). The behavior questions are more significant; students who carry credit card debt complete fewer questions sooner. Interestingly, students who sometimes pull "all-nighters" actually complete more questions sooner, which suggests they are aware of their tendency to procrastinate and want to overcome it.

One measurable real-life behavior we have is on which day subjects chose to complete the experiment: in order to make it more academically friendly, subjects could choose which day of the week was best for them, and complete the survey on that week. Table 5 shows the differences between subjects who completed it on

TABLE 3. Demographics

|  | $(1)$ | $(2)$ | $(3)$ | $(4)$ |
| :--- | :---: | :---: | :---: | :---: |
| OLS3 | Tobit3 | Interval3 | IntervalNew3 |  |
| Female | 1.259 | 1.331 | 1.270 | 1.259 |
|  | $(1.33)$ | $(1.22)$ | $(1.08)$ | $(1.14)$ |
| Other | 5.956 | 7.558 | 8.748 | 7.929 |
|  | $(1.28)$ | $(1.32)$ | $(1.25)$ | $(1.22)$ |
| Sophomore | $-2.715^{* * *}$ | $-3.370^{* * *}$ | $-3.550^{* * *}$ | $-3.275^{* * *}$ |
|  | $(-2.61)$ | $(-2.76)$ | $(-2.75)$ | $(-2.70)$ |
| Junior | -1.809 | -1.968 | -2.310 | -2.226 |
|  | $(-1.20)$ | $(-1.14)$ | $(-1.20)$ | $(-1.24)$ |
| Senior | $-5.382^{* *}$ | $-6.397^{* *}$ | $-6.720^{* *}$ | $-6.287^{* *}$ |
|  | $(-2.25)$ | $(-2.26)$ | $(-2.30)$ | $(-2.30)$ |
| Journalism | 2.371 | 2.715 | 2.866 | 2.758 |
|  | $(1.22)$ | $(1.22)$ | $(1.18)$ | $(1.20)$ |
| SocialScience | -0.399 | -0.560 | -0.782 | -0.650 |
|  | $(-0.20)$ | $(-0.25)$ | $(-0.32)$ | $(-0.28)$ |
| Other | 2.490 | 2.743 | 2.843 | 2.751 |
|  | $(1.47)$ | $(1.43)$ | $(1.36)$ | $(1.39)$ |
| Chinese | -0.607 | -0.798 | -0.807 | -0.736 |
|  | $(-0.50)$ | $(-0.57)$ | $(-0.53)$ | $(-0.52)$ |
| Arabic | 3.683 | 4.126 | 4.369 | 4.166 |
|  | $(1.24)$ | $(1.18)$ | $(1.17)$ | $(1.19)$ |
| Other | $-2.902^{* *}$ | $-3.085^{*}$ | $-3.535^{* *}$ | $-3.393^{* *}$ |
|  | $(-2.01)$ | $(-1.86)$ | $(-1.97)$ | $(-2.02)$ |
| Constant | $7.211^{* * *}$ | $5.500^{* *}$ | $8.804^{* * *}$ | $8.643^{* * *}$ |
|  | $(3.28)$ | $(2.19)$ | $(3.25)$ | $(3.37)$ |
| All Controls | yes | yes | yes | yes |
| N | 2619 | 2619 | 2619 | 2619 |
| $t$ statistics in parentheses |  |  |  |  |
| $* p<0.10,{ }^{* *} p<0.05,{ }^{* * *} p<0.01$ |  |  |  |  |

TABLE 4. Behaviors

|  | $(1)$ | $(2)$ | $(3)$ |  |
| :--- | :---: | :---: | :---: | :---: |
| OLS2 | Tobit2 | Interval2 | $(4)$ <br> IntervalNew2 |  |
| Delay | 0.326 | 0.370 | 0.374 | 0.358 |
| Unexpected | $(0.59)$ | $(0.56)$ | $(0.52)$ | $(0.54)$ |
|  | 0.0980 | 0.0861 | 0.0881 | 0.0936 |
| Schedule* | $(0.18)$ | $(0.14)$ | $(0.13)$ | $(0.15)$ |
|  | -0.108 | -0.188 | -0.157 | -0.130 |
| Sooner* | $(-0.19)$ | $(-0.28)$ | $(-0.22)$ | $(-0.20)$ |
|  | $1.317^{* *}$ | $1.562^{* *}$ | $1.492^{*}$ | $1.409^{*}$ |
| Accomplish* | $(2.00)$ | $(2.00)$ | $(1.76)$ | $(1.78)$ |
|  | 0.333 | 0.342 | 0.556 | 0.503 |
| Control | $(0.53)$ | $(0.46)$ | $(0.69)$ | $(0.67)$ |
|  | 0.656 | 0.711 | 0.811 | 0.764 |
| Comfortable | $(1.00)$ | $(0.92)$ | $(0.96)$ | $(0.97)$ |
|  | 0.0959 | 0.139 | 0.184 | 0.157 |
|  | $(0.16)$ | $(0.19)$ | $(0.23)$ | $(0.21)$ |
| CreditCard* | $-2.032^{* *}$ | $-2.448^{* *}$ | $-2.553^{* *}$ | $-2.386^{* *}$ |
|  | $(-2.20)$ | $(-2.26)$ | $(-2.17)$ | $(-2.16)$ |
| Bank* | -0.216 | -0.0958 | -0.171 | -0.194 |
|  | $(-0.16)$ | $(-0.06)$ | $(-0.10)$ | $(-0.12)$ |
| AllNighter* | $2.532^{* *}$ | $2.991^{* *}$ | $3.289^{* *}$ | $3.071^{* *}$ |
|  | $(2.46)$ | $(2.45)$ | $(2.47)$ | $(2.47)$ |
| Late* | -1.663 | -1.749 | -2.294 | -2.148 |
|  | $(-1.42)$ | $(-1.31)$ | $(-1.64)$ | $(-1.62)$ |
| Constant | 1.672 | -0.759 | 1.485 | 1.853 |
|  | $(0.42)$ | $(-0.16)$ | $(0.30)$ | $(0.40)$ |
| All Controls | yes | yes | yes | yes |
| N | 2628 | 2628 | 2628 | 2628 |

TABLE 5. Day of the Week

|  | (1) <br> WeekOLS | (2) <br> WeekTobit | (3) <br> WeekInterval | $\begin{gathered} \hline \hline(4) \\ \text { WeekInterval2 } \end{gathered}$ |
| :---: | :---: | :---: | :---: | :---: |
| Sunday | 0 | 0 | 0 | 0 |
|  | (.) | (.) | (.) | (.) |
| Monday | $\begin{aligned} & 1.634 \\ & (0.71) \end{aligned}$ | $\begin{aligned} & 1.873 \\ & (0.68) \end{aligned}$ | $\begin{aligned} & 2.092 \\ & (0.69) \end{aligned}$ | $\begin{aligned} & 1.976 \\ & (0.70) \end{aligned}$ |
| Tuesday | $\begin{aligned} & -1.725 \\ & (-1.00) \end{aligned}$ | $\begin{aligned} & -1.987 \\ & (-0.97) \end{aligned}$ | $\begin{aligned} & -2.212 \\ & (-0.98) \end{aligned}$ | $\begin{aligned} & -2.057 \\ & (-0.98) \end{aligned}$ |
| Wednesday | $\begin{aligned} & -1.023 \\ & (-0.58) \end{aligned}$ | $\begin{aligned} & -1.356 \\ & (-0.65) \end{aligned}$ | $\begin{aligned} & -1.727 \\ & (-0.78) \end{aligned}$ | $\begin{aligned} & -1.482 \\ & (-0.71) \end{aligned}$ |
| Thursday | $\begin{aligned} & -0.634 \\ & (-0.35) \end{aligned}$ | $\begin{aligned} & -0.878 \\ & (-0.41) \end{aligned}$ | $\begin{aligned} & -1.161 \\ & (-0.49) \end{aligned}$ | $\begin{aligned} & -1.005 \\ & (-0.45) \end{aligned}$ |
| Friday | $\begin{gathered} -5.864^{* * *} \\ (-2.96) \end{gathered}$ | $\begin{gathered} -6.664^{* * *} \\ (-2.90) \end{gathered}$ | $\begin{gathered} -7.333^{* * *} \\ (-2.97) \end{gathered}$ | $\begin{gathered} -6.888^{* * *} \\ (-2.96) \end{gathered}$ |
| Saturday | $\begin{gathered} -7.157^{* * *} \\ (-4.52) \end{gathered}$ | $\begin{gathered} -8.396^{* * *} \\ (-4.40) \end{gathered}$ | $\begin{gathered} -9.062^{* * *} \\ (-4.29) \end{gathered}$ | $\begin{gathered} -8.443^{* * *} \\ (-4.31) \end{gathered}$ |
| Late | $\begin{aligned} & -0.818 \\ & (-0.39) \end{aligned}$ | $\begin{aligned} & -1.338 \\ & (-0.55) \end{aligned}$ | $\begin{aligned} & -1.839 \\ & (-0.67) \end{aligned}$ | $\begin{aligned} & -1.522 \\ & (-0.59) \end{aligned}$ |
| All Controls N | 2684 | 2684 | 2684 | 2684 |

$t$ statistics in parentheses
${ }^{*} p<0.10,{ }^{* *} p<0.05,{ }^{* * *} p<0.01$
days of the week. Not surprisingly, subjects who put off completing the experiment until Friday or Saturday choose to do a great deal fewer questions in early time periods. In fact, these subjects may be more representative of a typical effort allocation decision-maker, who didn't specifically plan on making decisions on that particular day.

Overall, our reduced form results suggest that subjects respond more to differences in exchange rates, rather than differences in time. However, the

TABLE 6. Question Statistics

|  | Week 6 | Week 8 | Week 10 |
| :--- | :---: | :---: | :---: |
| Correct Answers | $71.7 \%$ | $74.0 \%$ | $71.1 \%$ |
| Completed Questions | $98.1 \%$ | $79.7 \%$ | $78.2 \%$ |
| Total Credit | $71.7 \%$ | $59.0 \%$ | $55.6 \%$ |

structure of these regressions only allows comparison between different allocations, rather than the results of individual allocations - subjects who put off questions in every allocation cannot drive estimates of the coefficients on time.

Also, our results do not account for the possibility of different perceived costs and benefits of questions in different time periods. It is possible that subjects anticipate that they may forget class material, making questions more difficult in future periods, or that they have a smaller likelihood of remembering to complete them. Aggregate results show such fears about forgetting may be appropriate. Table 6 shows the percentage of questions attempted and correctly answered in each time period: while the percentage of answers that were correct didn't appreciably change between time periods, many more questions were forgotten in the later time periods. It is possible that students anticipated this, but we cannot control for this in the reduced form. Using structural analysis, however, allows us to potentially control for different perceived cost in different time periods, and to use information within each decisions, rather than only between comparisons.

Given the imperfect follow-through of students, it may be possible that instead of expressing their true time preferences in this experiment, some subjects are attempting to use this experiment as a commitment device, forcing themselves to complete tasks sooner than they might otherwise prefer. The structure of the experiment makes this very possible: subjects knew exactly what they were getting into in the experiment (which was necessary given the academic environment), and thus were likely thinking of these decisions in a more "rational", long-term
mindset than in other experiments on time preferences. This finding suggests that the conditions of time preference experiments can affect the results.

## Structural Results

Given our utility assumptions, the following equation describes the optimal number of tasks to complete, when subjects are minimizing disutility of effort. In this estimation, we restrict $\alpha>1$, because this implies interior solutions to the budget decisions, which make up a large majority of observations.

$$
\begin{gather*}
x_{t}=\frac{\left(\frac{m}{P_{2}}+\omega_{t+k}\right)\left(\frac{\gamma_{t+k} P_{1}}{\gamma_{t} P_{2}}\right)^{\frac{1}{(\alpha-1)}}}{\left(1+\frac{P_{1}}{P_{2}}\left(\frac{\gamma_{t+k} P_{1}}{\gamma_{t} P_{2}}\right)^{\frac{1}{(\alpha-1)}}\right)}-\frac{\omega_{t}}{\left(1+\frac{P_{1}}{P_{2}}\left(\frac{\gamma_{t+k} P_{1}}{\gamma_{t} P_{2}}\right)^{\frac{1}{(\alpha-1)}}\right)}  \tag{3.12}\\
\omega_{t}=\lambda_{t}+\sum_{i=1}^{n} \theta_{i} * \eta_{i} \tag{3.13}
\end{gather*}
$$

Table 7 shows the results from estimating this equation using Nonlinear Least Squares estimation. Future work will investigate and implement the use of interval regressions for nonlinear applications; there is no pre-packaged tool for this. Different columns include different assumptions about the structure of $\lambda$ : columns 1 and 3 assume $\lambda$ is constant across time periods, while column 2 allows $\lambda$ to vary between the three different time periods that appeared in the experiment. Column 3 also includes multiplicative scaling of questions in periods 2 and 3: each week scales up the effort cost of questions by multiplier.

TABLE 7. NLS Structural Results

|  | $(1)$ <br> Default | $(2)$ <br> ChangingBaseline | $(3)$ <br> Multiplicative |
| :--- | :---: | :---: | :---: |
| beta | $1.044^{* * *}$ | $0.594^{* * *}$ | $0.632^{* * *}$ |
|  | $(0.142)$ | $(0.0539)$ | $(0.0916)$ |
| delta | $0.723^{* * *}$ | $1.022^{* * *}$ | $0.909^{* * *}$ |
|  | $(0.0473)$ | $(0.0302)$ | $(0.0281)$ |
| Alpha | $1.969^{* * *}$ | $1.271^{* * *}$ | $1.555^{* * *}$ |
|  | $(0.376)$ | $(0.108)$ | $(0.195)$ |
| lambda1 | $-32.52^{* * *}$ | $-29.59^{* * *}$ | $-30.58^{* * *}$ |
|  | $(2.451)$ | $(1.503)$ | $(1.863)$ |
| lambda2 |  | $-27.87^{* * *}$ |  |
|  |  | $(1.599)$ |  |
| lambda3 |  | $-20.98^{* * *}$ |  |
|  |  | $(1.666)$ |  |
| multiplier |  |  | $0.0967^{* * *}$ |
|  |  |  | $(0.0200)$ |
| work | 0.140 | -0.0228 | 0.0819 |
|  | $(0.548)$ | $(0.411)$ | $(0.495)$ |
| assign | $-1.390^{*}$ | $-1.279^{* *}$ | $-1.350^{* *}$ |
|  | $(0.552)$ | $(0.459)$ | $(0.519)$ |
| family | 0.199 | 0.0159 | 0.104 |
|  | $(0.458)$ | $(0.362)$ | $(0.423)$ |
| social | 0.503 | 0.499 | 0.515 |
|  | $(0.461)$ | $(0.412)$ | $(0.441)$ |
| $N$ | 2684 | 2684 | 2684 |
| adj. $R^{2}$ | 0.860 | 0.861 | 0.861 |
| Stard |  |  |  |

Standard errors in parentheses
${ }^{*} p<0.05,{ }^{* *} p<0.01,{ }^{* * *} p<0.001$

Of the categories of outside effort, only time spent on assignments is significant, and the coefficient is negative, suggesting subjects either substitute from assignment time toward review questions, or already anticipate spending time on review questions and count that as assignment time. However, we are using the full sample, which includes many patently unrealistic responses of outside effort (like more than 24 hours in a day), or responses that subjects were not confident about, suggesting they may not reflect reality. A more accurate measure of outside effort may present different results.

Given the qualitative results which seem to indicate completing tasks sooner than necessary, it is slightly surprising to see values of $\beta$ and $\delta$ that generally align with previous literature. When we control for different costs of effort across different time periods, we find that $\beta$ equals about .6 , statistically different from 1 , which is lower than other studies have found. ( $\delta$ is not statistically different from 1 under those specifications.) This shows that our results may not actually be driven by "preproperation", but instead by different mental accounting of questions in different time periods.

One of the most striking results from this estimation is that $\lambda$ appears negative, usually around -30 . This suggests that subjects are prepared to complete 30 questions in any particular time period, and only begin to (expect to) feel disutility after the first 30 questions. Interpreting this results for when subjects allocate fewer than 30 questions to a day is difficult - our utility function isn't even defined for negative values of $\omega+x$, except for very specific values of $\alpha$. Unfortunately, constraining $\lambda$ to be positive in regressions creates non-convergence.

Future work will determine the significance and accuracy of the finding of $\lambda<0$, as well as investigating the alternative perspective of increasing effort being coded as leisure loss.

## Loss of Leisure Specification

An alternative perspective to take on this whole subject is to view assigned questions as losses of leisure, instead of increased cost. Given this perspective, with diminishing marginal (positive) utility of leisure implying interior solutions instead of increasing marginal costs of effort (so we restrict $\alpha<1$ instead of $\alpha>1$ ), we find some very interesting results. Here, endowed effort appears positive $(\lambda>0)$, but its scale implies negative leisure for many observed numbers of questions. Table 8 shows results from this estimation.

Interestingly, given the losses treatment, we actually find significant evidence of "preproperation" - $\beta$ is significantly greater than 1 , which suggests subjects weigh future leisure with more weight than present leisure. This result is dramatically different than both our other perspective, and previous studies. The fact that we find stable results with dramatically different estimates for two different treatments suggests other published results may depend dramatically on the perspective taken about effort.

### 3.6 Discussion and Conclusion

This experiment on the discounting of effort over time finds little evidence of strong discounting behavior. While subjects choices are quite responsive to differences in the exchange rate between doing questions now rather than later, they are much less responsive to changes in the time delay. Some kinds of outside effort and commitments, as measured by self-reports, affect when subjects want to do the questions, but the coefficients are small. Results from a structural estimation with specific functional forms for utility and discounting vary depending

TABLE 8. Losses Perspective

|  | $(1)$ <br> Default2 | $(2)$ <br> ChangingBaseline2 | $(3)$ <br> Multiplicative2 |
| :--- | :---: | :---: | :---: |
| work | 0.233 | 0.139 | 0.109 |
| assign | $(0.530)$ | $(0.499)$ | $(0.406)$ |
|  | 0.632 | 0.589 | $0.686^{*}$ |
| classtime | $(0.430)$ | $(0.412)$ | $(0.386)$ |
|  | -0.378 | -0.422 | -0.330 |
| study | $(0.385)$ | $(0.360)$ | $(0.309)$ |
|  | 0.0453 | 0.0727 | 0.0324 |
| clubs | $(0.342)$ | $(0.338)$ | $(0.309)$ |
|  | 0.425 | 0.472 | 0.324 |
| family | $(0.499)$ | $(0.489)$ | $(0.445)$ |
|  | $1.333^{* * *}$ | $1.244^{* * *}$ | $1.012^{* *}$ |
|  | $(0.479)$ | $(0.470)$ | $(0.429)$ |
| social | $-0.788^{*}$ | $-0.696^{*}$ | $-0.893^{* *}$ |
|  | $(0.414)$ | $(0.404)$ | $(0.388)$ |
| beta | $1.114^{* * *}$ | $1.406^{* * *}$ | $1.505^{* * *}$ |
|  | $(0.102)$ | $(0.272)$ | $(0.458)$ |
| delta | $0.829^{* * *}$ | $0.827^{* * *}$ | $0.630^{* * *}$ |
|  | $(0.0428)$ | $(0.0628)$ | $(0.143)$ |
| Alpha | 0.136 | 0.328 | 0.301 |
|  | $(0.369)$ | $(0.318)$ | $(0.207)$ |
| Lambda | $18.35^{* * *}$ | $22.12^{* * *}$ | $17.97^{* * *}$ |
|  | $(2.807)$ | $(2.554)$ | $(1.661)$ |
| lambda8 |  | $19.05^{* * *}$ |  |
| lambda10 |  | $(2.732)$ |  |
|  |  | $19.53^{* * *}$ | $(2.174)$ |
| multiplier |  |  | $-0.0715^{* * *}$ |
|  |  |  | $(0.0226)$ |
| Observations | 2684 | 2684 | 2684 |
|  |  |  |  |
|  |  |  |  |

Standard errors in parentheses
${ }^{*} p<0.10,{ }^{* *} p<0.05,{ }^{* * *} p<0.01$
on the assumptions about mental accounting of effort - is it a loss of the good of leisure, or an increase in the bad of effort?

One possible explanation for these mixed results, which was not considered in the design, is that some subjects may have been using the experiment as a commitment device (which sometimes failed), and that the estimates are therefore a mix of preferences and efforts to commit to choices for which preferences may change over time. This may explain some coefficient and parameter estimates, and aligns with some subjects' failure to follow through with their allocations choices. Therefore, while our results for time preferences significantly differ from the most related work, Augenblick et al. (2015) and Augenblick and Rabin (aper), this may be due to differences in the circumstances surrounding the experiment.

This work shows that the design and framing of time allocation experiments can make a large difference in how subjects treat the experiment. When subjects are given significant warning and preparation time about real-effort tasks, they make very different decisions than if they enter with few expectations. Running an experiment in an academic setting limited our ability to suppress subjects' expectations, and thus our ability to truly elicit time preferences. Work in similar academic settings (such as Burger et al. (2011)) was not limited by the use of graded credit as a motivation, but found similar behavior in attempting to implement self-control.

We have also shown that the amount of effort, outside the experiment, that subjects expected to have to do is worth considering when investigating time allocation decisions. Previous work has either ignored this, or dismissed it as a confound that cannot be investigated. In future work I will investigate the decisionmaking of individuals and heterogeneity in their decisions.

While mental accounting of money has been extensively studied (especially through Prospect Theory), significantly less research has investigated mental accounting of effort. Effort can be framed as an increase in a bad, or as a loss of a good, i.e. leisure. Many researchers have found differences between mental accounting of losses versus gains for money (including Estle et al. (2006) and Appelt et al. (2011)), so finding similar results for effort is not surprising. Future planned experiments will investigate the application of Prospect Theory and mental accounting to effort versus leisure, and compare that to preferences over money.

In conclusion, this experiment has shown that experimentally elicited time preferences for effort are not the same as the time preferences of money that are typically found in experiments. Additionally, while outside effort as measured by our elicitation method did not often show up as statistically or economically significant, among some samples of subjects it may make a difference and is worth increased investigation in the future.

## CHAPTER IV

## RISK PREFERENCES FOR REAL EFFORT TASKS

Risk preferences are fundamental to any model of economic behavior. While risk preferences, specifically utility curvature and probability weighting, have been extensively studied for money, very little empirical research has been completed on risk preferences for things other than money. We complete an experiment with two separate sessions: one which elicits risk preferences for money, the other which elicits risk preferences for leisure and effort, defined as the quantity of real-effort tasks to complete. In aggregate, subjects were more risk-neutral over the quantity of leisure compared to money (less risk-averse over gains, less risk-seeking over losses). Individual risk preferences for money versus leisure were not correlated, suggesting individuals have many perspectives in evaluating risk.

### 4.1 Introduction

Real-effort tasks are becoming more and more popular in economics experiments, as a way to induce disutility in subjects, to measure time preferences, to promote a sense of ownership of payoffs, and many other possibilities. However, there is only a limited understanding of the relationship between these tasks and (dis)utility, especially when connected to risk. Specifically, we don't know the connection between risk preferences for effort and leisure versus the more well-studied risk preferences for money. Individuals could have consistent risk preferences across various domains, being risk-averse or risk-seeking for many different outcomes, or these preferences could be independent within individuals. If risk preferences are consistent across domains, then research insights gained from
studying decisions over money can be generalized to other outcomes. However, if there is little connection between preferences over different domains, then we cannot carry over previous results into different domains, and further research is necessary. Specifically, if we continue to use real-effort tasks in experiments, we as economists need to better understand the utility and preferences behind these real-effort tasks, instead of assuming they are comparable with money.

In a novel experimental treatment that elicited risk preferences for both money and leisure/effort across two different sessions, we find that these preferences are independent - while similar on an aggregate level, decisions over money and decisions over leisure are correlated with different covariates, and at an individual level only slightly correlated with each other. We conclude that economists need to study risk preferences in domains other than money, and not assume equal risk preferences when comparing outcomes in different domains, such as money versus time or health or direct consumption.

Section 2 presents a literature review of the basics of risk and real-effort tasks; section 3 describes the experiment, including two separate sessions: one for money, and one for effort. Section 4 shows the utility model and empirical strategies used in estimation, section 5 presents the aggregate results of those estimations, while section 6 includes results of individual estimations and their connections. Section 6 concludes.

### 4.2 Literature Review

Like countless other papers before it, this experiment aims to estimate risk attitudes of individuals. Risk is fundamental to nearly all fields of economics, and knowing how individuals regard risk has been a considerable goal throughout much of the history of experimental economics. Yet few have analyzed risk preferences as
they relate to anything other than money, including consumption goods, leisure, or in our case, real-effort tasks. In this paper, we extend the literature by looking at the influence of real-effort tasks on risk preferences.

Harrison and Rutström (2008) reviews a wide range of elicitation and estimation methods for risk attitudes, including many common but imprecise techniques, and others with more economic precision. One approach is the Multiple Price List (MPL), popularized by Holt and Laury (2002), which has been used in many other successful contexts (including Harrison et al. (2007)). Another technique is a simple lottery comparison method known as Random Lottery Pairs (RLP) originated by Hey and Orme (1994) and used in many applications, including Harrison and Rutström (2009), where subjects are shown two lotteries and have to choose between them, with no systematic presentation. While the lottery comparison technique does not allow as much precision in calculating the parameters of risk preferences, it allows a wider variety of comparisons, including comparisons using more than two outcomes, and a broader range of estimation techniques. It is also very transparent and easy for experimental subjects to understand. For these reasons, and for comparing results with previous experiments, we use the lottery comparison technique.

The dominant theory of decision-making under risk has been Expected Utility Theory. However, there are many known deficiencies in this theory: specifically, observable behavior that cannot be explained by this theory, such as the Allais paradox. Many attempts have been made to explain apparent anomalies in risk preferences and deviations from expected utility theory, including prospect theory with loss aversion (Kahneman and Tversky (1979) and Tversky and Kahneman (1991)) and probability weighting (Prelec and Loewenstein (1998)). These predict a four-fold pattern of risk seeking and aversion over small and large probability gains
and losses: see Harbaugh et al. (2002) and Harbaugh et al. (2010) for a discussion of this pattern. Harrison and Rutström (2009) discusses some of these approaches, finding that given their experimental results, a mixture of expected utility and prospect theory best describes observed choices. We aim to investigate whether this four-fold pattern applies to both money and effort.

Real-effort tasks, on the other hand, occur sporadically throughout the literature in various topics, including related to time preferences (Augenblick et al. (2015)), procrastination (Bisin and Hyndman (2014a)), and especially competition (examples include Niederle and Vesterlund (2007) and Dreber et al. (2014)). There is an even larger literature of effort allocation without real-effort tasks, especially in competition, where subjects pick a number to plug into a cost function (such as Harbring and Irlenbusch (2005)), but real-effort tasks are a completely different domain in a subject's decision-making, and thus results from these studies are not very useful in identifying separate risk preferences. While some real-effort tasks (like mathematical problems often used in competition studies) demonstrate significant differences in skill and willingness to compete by demographics, studies investigating the effects of real-effort tasks usually use tasks with little to no skill or learning effects, little variation in difficulty between subjects, and which provide no obvious benefit to anyone. A prime example of this is the digit-counting task from Abeler et al. (2011), which we use in our experiment.

Abeler et al. (2011) in fact is one of the few papers that deals with effort given risky rewards. While they introduce a theoretical framework regarding cost of effort and utility maximization, they do not attempt to calculate cost of effort and thus the degree of risk aversion in subjects. On the other hand, Gill and Prowse (2012) does calculate the cost of effort, but assumes it takes quadratic form, and because the given experiment involves competition, the probabilities
that subjects face are not explicit, making calculation of risk aversion difficult. Similarly, in Carpenter et al. (2007), subjects make real-effort allocation decisions in a competitive environment, and costs of effort are calculated, but with restrictive assumptions and unknown probabilities of outcomes.

Different risk preferences in different domains have received some attention; outside of economics, the DOSPERT (DOmain-SPEcific Risk Taking) survey (Blais and Weber (2006)) attempts to measure five categories of risk-taking and perceptions, but is limited by being only a survey. Within economics, differences in risk preferences for the domains of losses versus gains is one of the primary components of prospect theory, and many have performed experiments investigating this (including Harrison and Rutström (2009)). However, in the domain of effort and leisure, labeling losses and gains is not as easy as positive or negative numbers: a positive number of additional tasks to complete is a loss of leisure, which we assume is a good. Most research using real-effort tasks (including Abeler et al. (2011) and Augenblick et al. (2015)) analyze only the disutility of additional effort, not the gained utility from increased leisure. While we use both approaches in our analysis, we prefer the perspective of leisure, which as positive good is more directly comparable to money.

### 4.3 Experimental Design

Subjects were recruited from sections of introductory classes in economics, with no formal introduction to risk preferences (although they were likely to have seen mathematical calculations of expected value). Subjects signed up for two different experiment sessions, one session with payoffs of real-effort tasks to complete, the other with outcomes of monetary payoffs, without knowing which session was which beforehand. The two sessions were 1-3 weeks apart, depending on
individual schedules. Sessions consisted of up to 12 individuals (mean subjects per session $=6.1$ ). Each session, completed using the z-Tree software (from Fischbacher (2007)) included presentations of both gains and losses. For money, gains were labeled as earning additional money from a $\$ 0$ baseline, while losses were labeled as losing money from a $\$ 50$ baseline. For the leisure gains/effort loss frame, subjects were instructed that they would be completing 25 real-effort tasks at the end of the experiment, and lotteries entered were for relief from these 25 tasks - subjects would gain leisure from having to complete fewer tasks. In the leisure loss/effort gain frame for the real-effort tasks sessions, the default was zero tasks completed, and lotteries were over having to complete additional tasks. Sections were randomly assigned to begin with either the gains or losses frame, and switch to the other frame halfway through the experiment.

The real-effort task involved in this research was taken from Abeler et al. (2011): subjects were shown an image of a $15 \times 15$ matrix of digits ( 225 total digits), consisting only of 0 s and 1 s . They had to count the total number of 0 s appearing in the matrix and input this total number. If they inputted the correct answer or a number within two of the correct answer, the task was completed; if wrong, they had one more chance. If wrong again, the task counted as a negative completion and they were shown a new table - this was to discourage random guessing. See Figure 5 for an example. This was designed to be boring, taxing concentration, and with obviously no outside benefit to anyone, thus giving us a task all subjects would try to avoid. There was also only a small possibility that subjects would substantially improve in their ability to complete this task over the time period of this experiment, although some subjects developed new techniques for digitcounting and did marginally improve their speed.

FIGURE 5. Real Effort Task


In the sessions involving effort, subjects first completed as many tasks as they could in five minutes while paid a piece rate for these tasks; this simultaneously informed them of the difficulty and effort required in the task, their ability to complete these tasks, and set a baseline of the task completion rate for each subject. Once they were sufficiently familiar with these tasks, subjects were shown 60 comparisons of lotteries through z-Tree software in two sets of 30 comparisons, and had to choose which one they would prefer to enter. The lotteries had possible outcomes of $5,10,15$, and 20 tasks to complete, although each lottery comparison only included three possible prizes. Probabilities were in multiples of $1 / 8$, including a few degenerate lotteries with certain payoffs. Comparisons were presented as two pie charts, accompanied by tables of outcomes and probabilities.

The first and second sets of 30 comparisons were in opposite framings one was shown as additional tasks to complete, the other as relief from 25 tasks. Before each of these sets of comparisons, subjects took a short comprehension quiz, both to make sure subjects understood the experiment and to preview the different framings. Figure 6 shows the same lottery comparison, presented in the two different framings. Notice that we use a different color palette for each framing, to minimize association in decision-making between the two frames. After answering demographic questions and a short numeracy quiz, one of the 60 lottery comparisons was randomly chosen to apply to each subject. Then, the lottery that the subject selected from that comparison was resolved, and subjects were assigned a number of tasks to complete. At the end of the experiment, subjects had to complete however many real-effort tasks they were assigned. As soon as they finished, they collected their payments (a show-up fee and a lump sum payment for completing tasks) and were free to leave. Subjects could leave early before completing all of their assigned tasks, but were not given the completion payment
(only two subjects left early, after making an effort to complete their tasks). Note that the number of tasks assigned did not affect the size of the completion payment. Therefore, there was no benefit to having more tasks to complete, making number of tasks to complete (which we label as effort) a pure "bad", and being relieved of tasks (labeled leisure) a pure good.

In the sessions involving money, subjects also had to complete 60 total comparisons, in two rounds of 30 comparisons each, with outcomes of $\$ 10, \$ 20, \$ 30$, and $\$ 40$. Again, one of the rounds was presented as gains, this time as additional money, while the other was framed as losses from a baseline of $\$ 50$, with the order of the framings varying by session. Figure 7 shows the two framings of money comparisons. In both the money and effort sessions, subjects had eight minutes to complete each round of 30 comparisons, and could not move on in the experiment until the eight minutes were over. Therefore, there was no benefit to subjects completing comparisons faster than the given eight minutes. (There were penalties for not completing comparisons in 8 minutes that never had to be applied). After completing all of the comparisons, subjects completed a demographic survey, the DOSPERT assessment of risk attitudes, and a personality test. Once all students had completed all assessments, one of the 60 comparisons was chosen for each student, then the chosen lottery of that comparison was resolved, and subjects were paid in cash and dismissed.

Sessions occurred in January-March 2016, with subjects recruited from introlevel sections of economics courses. We collected data from 89 total subjects in sessions over money and 92 subjects in effort. Of those subjects, 79 completed both sessions: our analysis only uses these 79 subjects.

Demographics of the 79 subject sample are shown in Table 9. Note that there were slight discrepancies in reported major and class year between the two sessions

FIGURE 6. Effort Comparisons


FIGURE 7. Money Comparisons


TABLE 9. Demographics

|  | freq | pct |
| :--- | :---: | :---: |
| Male | 46 | 58.23 |
| Female | 33 | 41.77 |
| Economics | 13 | 16.46 |
| Business | 42 | 53.16 |
| Journalism | 10 | 12.66 |
| Social Science | 7 | 8.86 |
| Natural Science | 2 | 2.53 |
| Other | 5 | 6.33 |
| English | 65 | 82.28 |
| Chinese | 8 | 10.13 |
| Other | 6 | 7.59 |
| Freshman | 39 | 49.37 |
| Sophomore | 22 | 27.85 |
| Junior | 10 | 12.66 |
| Senior | 6 | 7.59 |
| $N$ | 79 |  |

(and two missing class years); we use reported demographics from the money session.

### 4.4 Test Strategy

Within each comparison, we assume that subjects select the lottery with the greater utility. We have multiple models of utility, all applicable to money and effort/leisure.

The first model is a simple Constant Relative Risk Aversion utility function, where utility is evaluated based on final outcomes, not subject to framing effects:

$$
u(x)=\frac{\left(x^{\alpha}\right)-1}{\alpha}
$$

In this model, $\alpha$ is the curvature of utility, and our measure of risk aversion. Other papers replace $\alpha$ with $1-\rho$ as a convention, but for familiarity, we use
$\alpha$. Subjects are risk-averse when $\alpha<1$, risk-neutral when $\alpha=1$, and riskseeking when $\alpha>1$. However, this basic model assumes that the utility of gaining $\$ 30$ (compared to gaining $\$ 0$ ) is the same as losing $\$ 20$ from a baseline of $\$ 50$ (compared to losing all $\$ 50$ ), which is not consistent with observed behavior.

The other model of utility, based on Prospect Theory, assumes a different utility curvature for losses. Here, outcomes are in comparison to a reference point, and thus subject to framing effects:

$$
u(x)= \begin{cases}\frac{\left(x^{\alpha}\right)-1}{\alpha} & \text { if } x \geq 0 \\ -\frac{\left(-x^{\beta}\right)-1}{\beta} & \text { if } x<0\end{cases}
$$

Including both $\alpha$ and $\beta$ is described as the two-parameter model.
For gains, risk aversion and risk seeking are again marked by $\alpha<1$ and $\alpha>1$, respectively. However, $\beta$ marks the curvature of a negative function, and thus $\beta<1$ denotes risk-seeking and $\beta>1$ shows risk aversion. While $\alpha<1$ means that low-probability large gains are less attractive than they would be under risk neutrality and thus not worth seeking out, $\beta<1$ means low-probability large losses are not as prohibitive as they would be under risk neutrality, and thus not worth avoiding. Previous experiments in Prospect Theory included in the literature review have found $\alpha<1$ and $\beta<1$ - subjects are risk-averse in gains, and riskseeking in losses.

Note that we do not include utility from the reference point (either $\$ 50$ or 25 tasks) in these utility functions. Because we are generally assuming that decisions are based on utility differences of the two lotteries in the comparison, any term that is equal between the two lotteries would be canceled out by taking the difference. Therefore, reference-point utility is canceled out, and we cannot estimate it.

Traditionally, Prospect Theory includes a scaling term of $\lambda$ applied to losses, which reflects that the disutility of losses is greater in magnitude than the utility of same-sized gains - that is, a kink in the value function at zero. However, because our experiment never includes comparisons of gains against losses (only gains against gains and losses against losses), our data cannot identify a $\lambda$ term, so we do not include it in our utility model.

We can think of an increase in the number of tasks to complete as either an increase in a bad (effort) or a decrease of a good (leisure). Similarly, a decrease in tasks can be a decrease of effort or an increase in leisure. Because we primarily focus on utility of goods, rather than disutility of bads, we emphasize preferences for leisure instead of effort in our analysis. But in utility frameworks with gains and losses, $\alpha$ for leisure is the same as $\beta$ for effort.

Both $\alpha$ and $\beta$ can vary based on many different covariates. We will test whether they depend on demographics, differences in experimental procedure, risk attitudes as measured by the DOSPERT assessment, and measured Big 5 personality traits.

Subjects also have utilities of the lottery as a whole, as determined by the utilities of the possible outcomes and the probabilities of those outcomes. We have two models for this aggregation of utility. First is simple expected utility, where outcome utilities are weighted by their true probabilities (denoted $p(x)$ for each outcome $x$ ):

$$
U=\sum_{i=1}^{4} p\left(x_{i}\right) * u\left(x_{i}\right)
$$

Prospect Theory also introduces the possibility of probability weighting: subjects act as if they misperceive probabilities and use them to underweight or
overweight particular outcomes. We assume this weighting takes place through a function $w(p(x))$, where $p(x)$ is the probabilities as before.

$$
U=\sum_{i=1}^{4} w\left(p\left(x_{i}\right)\right) * u\left(x_{i}\right)
$$

The probability weighting function could have many possible functional forms, but we focus on the original functional form proposed by Kahneman and Tversky (1992):

$$
w\left(p_{k}\right)=p_{k}^{\gamma} /\left(\sum_{i=1}^{4} p_{i}^{\gamma}\right)^{1 / \gamma}
$$

In this functional form, a single parameter $\gamma$ determines the weighting of probabilities. Note that when $\gamma<1$, low probabilities are over-weighted and high probabilities are under-weighted, while when $\gamma>1$ the opposite is true. Most previous research has found overweighting of small probabilities $(\gamma<1)$.

Including $\gamma$ along with $\alpha$ and $\beta$ is described as the three-parameter model.
We assume that individuals then compare the utility of the two lotteries in the comparison, and select the lottery with higher utility. However, we assume that there is an additive error term $\epsilon$ in this decision-making process with a logistic distribution with scale parameter $\sigma$, proportional to the standard deviation. So, the estimated difference in utilities between lotteries A and B would be:

$$
\Delta U=U(A)-U(B)+\epsilon
$$

The actual probability of selecting a lottery, given a difference in utilities, follows the inverse logistic distribution. Consequently, the likelihood of observing Left or Right given a particular utility function also follows the inverse logistic function:

$$
P(A)=\exp \left(\frac{(U(A)-U(B))}{\sigma}\right) /\left(1+\exp \left(\left(\frac{(U(A)-U(B))}{\sigma}\right)\right)\right)
$$

Some previous work has ignored the importance of the scale parameter of the logistic distribution. Not including $\sigma$ would assume that the error-generating process has a scale of 1 in terms of absolute utility difference, a clearly overlystrong assumption, considering the wide variety of scales we may see in utility.

Considering that gains and losses are different domains with differing levels of familiarity and risk aversion, it is also possible that the error-generating process is different for gains versus losses, in which case we would estimate two separate $\sigma$ parameters, one for gains and one for losses.

Given this likelihood function, we use Maximum Likelihood estimation on the logarithm of this function. The parameters that result in the greatest likelihood function are the best estimate of the true parameters of the model.

Another possible assumption about the decision rule and the error-generating process is subjects compare ratios of utility (instead of absolute differences) which have multiplicative instead of additive errors. In this case, assuming positive utility:

$$
\text { if } \begin{cases}|U(A)| /|U(B)| * \epsilon>1 & \rightarrow \text { select A } \\ |U(A)| /|U(B)| * \epsilon<1 & \rightarrow \text { select B }\end{cases}
$$

The decision rule would be the opposite for negatively-valued utility, such as for losses of money.

Here, $\epsilon$ has a log-logistic distribution, also known as the Fisk distribution (such that $\log (\epsilon)$ has a logistic distribution) with median 1 and a scale parameter $\sigma$ (such that $\log (\epsilon)$ has median 0 and scale $\sigma$ as well).

Now let us transform this decision-making process:

$$
\begin{gathered}
|U(A)| /|U(B)| * \epsilon \gtreqless 1 \\
\log (|U(A)|)-\log (|U(B)|)+\log (\epsilon) \gtreqless 0
\end{gathered}
$$

We know $\log (\epsilon)$ has a logistic distribution from our previous assumptions. This looks very much like our previous decision rule for additive errors. So, let us similarly feed it into the inverse logistic function, for positive utility values:

$$
\begin{gathered}
P(A)=\exp \left(\frac{\log (U(A))-\log (U(B))}{\sigma}\right) /\left(1+\exp \left(\left(\frac{\log (U(A))-\log (U(B))}{\sigma}\right)\right)\right) \\
P(A)=\exp (\log (U(A)))^{1 / \sigma} /\left(\exp (\log (U(A)))^{1 / \sigma}+\exp (\log (U(B)))^{1 / \sigma}\right) \\
P(A)=\frac{U(A)^{1 / \sigma}}{U(A)^{1 / \sigma}+U(B)^{1 / \sigma}}
\end{gathered}
$$

This transformed probability is the scaled version of the Luce model (Luce (1959)). The traditional Luce specification with no exponential term assumes $\sigma=$ 1 , as we saw before.

For negative utility values, the decision rule runs the opposite direction, so the inverse logistic function gives probabilities of finding the alternative, and the actual probability is:

$$
P(A \mid U(A)<0)=1-\frac{|U(A)|^{1 / \sigma}}{|U(A)|^{1 / \sigma}+|U(B)|^{1 / \sigma}}=\frac{U(B)^{1 / \sigma}}{U(A)^{1 / \sigma}+U(B)^{1 / \sigma}}
$$

We will estimate utility parameters given this assumption about the decisionmaking and error-generating processes as well, and compare them to the results from additive errors.

### 4.5 Aggregate Results

The goal of our estimation is to find risk preferences for real-effort tasks, compared with risk preferences for money, for both gains and losses. We also want to investigate the role of covariates, such as demographics and personality traits, on those risk preferences. Finally, we wish to investigate the distribution of risk preferences for money and real-effort tasks among the subject population, and determine the relationship between risk preferences for money and effort in subjects. In order to achieve these goals, we also need to determine the decisionmaking process and accompanying error-generating process that best matches the observed choice patterns.

To determine the most appropriate error-generating process to assume, we take the two-parameter model (including $\alpha$ and $\beta$, but no $\gamma$ ) and run estimations
under all four of our possible stochastic assumptions: additive logistic errors with $\sigma=1$ (as occurs in many previous papers), $\sigma$ estimated, and two separate values of $\sigma$ estimated for gains and losses; also, multiplicative errors with $\sigma$ estimated. We then compare the predicted probabilities of choices under the model to the observed distribution of choices in the data. Table 10 shows the true distributions of choices for select lotteries in each domain (under Mean), and the distributions predicted for those lotteries, given various error structures. It also shows the mean difference in observed versus predicted distributions for all 60 lotteries (as Mean Abs. Error), given the two-parameter model. From this, we observe that for money, the Luce error structure performs slightly better than the $\sigma_{\text {Gain }} \neq \sigma_{\text {Loss }}$ model, both of which are better than the other two models. For effort, $\sigma_{\text {Gain }} \neq \sigma_{\text {Loss }}$ is by far the best model. Especially because of this much-improved performance in the effort domain, we use the $\sigma_{\text {Gain }} \neq \sigma_{\text {Loss }}$ model throughout the rest of the analysis.

TABLE 10. Error Structure Assumptions

| Money |  |  |  |  |  |  |
| :---: | :---: | :---: | ---: | ---: | ---: | ---: |
| Lottery | EV Diff | Mean | $\sigma=1$ | $\sigma_{\text {Gain }}=\sigma_{\text {Loss }}$ | $\sigma_{\text {Gain }} \neq \sigma_{\text {Loss }}$ | Luce |
| 4 | -3.75 | 0.4304 | 0.1789 | 0.2493 | 0.2557 | 0.3042 |
| 17 | 0 | 0.3544 | 0.5021 | 0.4872 | 0.4289 | 0.4170 |
| 38 | 1.25 | 0.7342 | 0.6269 | 0.5850 | 0.5698 | 0.5906 |
| 50 | 1.75 | 0.4557 | 0.6296 | 0.5911 | 0.5671 | 0.5519 |
| Mean Abs. Error |  |  |  |  |  |  |
| Effort |  | 0.1074 | 0.1043 | 0.0966 | 0.0915 |  |
| Lottery | EV Diff | Mean | $\sigma=1$ | $\sigma_{\text {Gain }}=\sigma_{\text {Loss }}$ | $\sigma_{\text {Gain }} \neq \sigma_{\text {Loss }}$ | Luce |
| 8 | 0 | 0.8228 | 0.4984 | 0.5021 | 0.6828 | 0.5239 |
| 15 | -1.875 | 0.8228 | 0.7975 | 0.7729 | 0.8001 | 0.7863 |
| 25 | 1.875 | 0.1519 | 0.1991 | 0.2161 | 0.1916 | 0.2090 |
| 49 | 1.25 | 0.2308 | 0.2874 | 0.3235 | 0.3417 | 0.3158 |
| Mean Abs. Error |  |  | 0.1225 | 0.1201 | 0.0694 | 0.1144 |

Given this error structure, we estimate the baseline parameters of the utility function in all three domains: money, effort, and leisure. First, money: Table

TABLE 11. Parameter Estimates for Money

|  | $(1)$ <br> choicesmoney | $(2)$ <br> choicesmoney | $(3)$ <br> choicesmoney |
| :--- | :---: | :---: | :---: |
| alpha | $0.435^{* *}$ | 0.353 | 0.430 |
|  | $(0.167)$ | $(0.230)$ | $(0.278)$ |
| beta |  | $1.367^{* * *}$ | $1.319^{* * *}$ |
|  |  | $(0.205)$ | $(0.208)$ |
| gamma |  |  | $1.047^{* * *}$ |
|  |  |  | $(0.0728)$ |
| sigmagain | 0.441 | 0.391 | 0.522 |
|  | $(0.245)$ | $(0.298)$ | $(0.501)$ |
| sigmaloss |  | 7.917 | 6.920 |
|  |  | $(4.965)$ | $(4.388)$ |
| $N$ | 4738 | 4738 | 4738 |
| AIC | 6025.8 | 6028.4 | 6028.6 |
| ll | -3010.9 | -3010.2 | -3009.3 |

Standard errors in parentheses
${ }^{*} p<0.05,{ }^{* *} p<0.01,{ }^{* * *} p<0.001$

11 shows parameter estimates for three different utility models: simple expected utility, expected utility with loss aversion, and probability-weighted utility with loss aversion. Standard errors in all regressions presented are clustered at the individual level. Note that small coding errors caused a few subjects to miss a single lottery comparison, so N is not divisible by 60 .

These results show that in all models, subjects are significantly risk-averse in the gains domain: $\alpha<1$. Subjects are also risk-averse in the losses domain: $\beta>1$. This finding of risk aversion in the losses domain runs counter to most previous research, which finds risk-seeking in losses. However, notice the Akaiki Information Criterion and the log-likelihood for each regression are very similar: controlling for loss aversion added very little to the predictive power, and actually caused an increase in the AIC. This shows that the difference in framing actually made very little difference in decision-making for subjects, assuming the framing
made losses induce negative utility. It is possible that instead of viewing losses as producing negative utility, subjects still view results as gains of money, just in a different domain and thus with a different risk aversion parameter. Table 14 shows results of this specification ("losses" are still positive, but with a different risk parameter) with the two-parameter for all three domains. But even though there is little evidence for the two-parameter model in aggregate, there may be differences in how outside factors affect these two parameters, so we continue to use this model in subsequent analysis.

For the three-parameter model, we find $\gamma$ very close to (and statistically indistinguishable from) one, and with minimal gain in accuracy. We can conclude that this experiment shows no evidence of probability weighting, and thus ignore the three-parameter model for future estimation.

Table 12 shows estimates for of parameters in preferences for leisure in the same three models as Table 11. First, we can dismiss the three-parameter model, as $\gamma$ is almost exactly equal to 1 . However, here we find that subjects are extremely risk-averse in both gains and losses of leisure: $\alpha$ is in a range characterized by Holt and Laury (2002) as "stay in bed", beyond "highly risk-averse". Inspection reveals that the main driver of this extreme risk aversion is avoiding the possibility of completing 20 tasks. Also, separating gains and losses makes a considerable improvement in estimation accuracy - there is a real difference in decisions over gains versus over losses. Table 14 also shows results for leisure with separating gains versus losses into two preference domains, instead of positive versus negative outcomes.

As an alternative framing for leisure, we can look at disutility over the number of tasks to complete, which we label as effort. Table 13 shows these estimates. Here, utility over increases in effort is negative, while utility from

TABLE 12. Parameter Estimates for Leisure

|  | $(1)$ <br> choicesleisure | $(2)$ <br> choicesleisure | $(3)$ <br> choicesleisure |
| :--- | :---: | :---: | :---: |
| alpha | $-0.783^{* * *}$ | $-1.311^{* * *}$ | $-1.310^{* * *}$ |
|  | $(0.213)$ | $(0.319)$ | $(0.326)$ |
| beta |  | $2.799^{* * *}$ | $2.796^{* * *}$ |
|  |  | $(0.393)$ | $(0.396)$ |
| gamma |  |  | $0.997^{* * *}$ |
|  |  |  | $(0.0485)$ |
| sigmagain | $0.0351^{* *}$ | 0.00962 | 0.00964 |
|  | $(0.0131)$ | $(0.00533)$ | $(0.00542)$ |
| sigmaloss |  | 272.8 | 270.7 |
|  |  | $(317.3)$ | $(318.4)$ |
| $N$ | 4722 | 4722 | 4722 |
| $A I C$ | 5562.9 | 5491.6 | 5493.6 |
| ll | -2779.4 | -2741.8 | -2741.8 |

Standard errors in parentheses
${ }^{*} p<0.05,{ }^{* *} p<0.01,{ }^{* * *} p<0.001$
decreases in effort is positive. The two- and three-parameter models are just the reverse of the models for leisure, and the parameter estimates reflect that: a loss of leisure is identical to a gain of effort. Naturally, these estimates again show extreme risk aversion.

Table 14 shows the effects of separating gains and losses, but not treating losses as negative. Note that in this reframing, $\alpha$ and $\beta$ for money are not statistically significantly different: there is no strong evidence for any loss aversion for money in our experiment. However, the regression is a better fit than the previous two-parameter model, suggesting that effort and leisure do still demonstrate a significant difference between gains and losses: subjects are more risk-averse in increases in leisure/decreases in effort. Comparing the AIC of all effort or leisure regressions, the best of all our specifications is the separated but not reframed effort regression: the decision-making process that best aligns with

TABLE 13. Parameter Estimates for Effort

|  | $(1)$ <br> choiceseffort | $(2)$ <br> choiceseffort | $(3)$ <br> choiceseffort |
| :--- | :---: | :---: | :---: |
| alpha | $3.832^{* * *}$ | $2.799^{* * *}$ | $2.796^{* * *}$ |
|  | $(0.352)$ | $(0.393)$ | $(0.396)$ |
| beta |  | $-1.311^{* * *}$ | $-1.310^{* * *}$ |
|  |  | $(0.319)$ | $(0.326)$ |
| gamma |  |  | $0.997^{* * *}$ |
|  |  |  | $(0.0485)$ |
| sigmagain | 4042.1 | 272.8 | 270.7 |
|  | $(4303.0)$ | $(317.6)$ | $(318.4)$ |
| sigmaloss |  | 0.00962 | 0.00964 |
|  |  | $(0.00533)$ | $(0.00543)$ |
| $N$ | 4722 | 4722 | 4722 |
| AIC | 5550.3 | 5491.6 | 5493.6 |
| ll | -2773.2 | -2741.8 | -2741.8 |

Standard errors in parentheses
${ }^{*} p<0.05,{ }^{* *} p<0.01,{ }^{* * *} p<0.001$
our data is that subjects minimize their disutility of effort, and do not view effort "losses" as such, but instead an alternative way of presenting increases in effort.

A reasonable question is whether order effects made a difference. While the order of the domains and of the framings was effectively randomized, it is still possible that they made a difference, even though they would not affect our aggregate estimates. Table 15 shows order effects: we find that subjects who attended the effort session second were even more risk averse over leisure gains, and subjects who saw money losses first before money gains were more risk-tolerant of losses. These small differences are not too surprising, and again do not matter for our aggregate results.

We also investigated if demographics made a difference in risk preferences. Table 16 shows the effects of demographics on $\alpha$ and $\beta$ in the two-parameter model (framed) for both money and leisure. Generally, major, language, and class year

TABLE 14. Separated Regressions, Not Reframed

|  | $(1)$ <br> choicesmoney | $(2)$ <br> choiceseffort | $(3)$ <br> choicesleisure |
| :--- | :---: | :---: | :---: |
| alpha | 0.353 | $2.799^{* * *}$ | $-1.041^{* * *}$ |
|  | $(0.230)$ | $(0.392)$ | $(0.249)$ |
| beta | $0.512^{* *}$ | $4.802^{* * *}$ | 0.00762 |
|  | $(0.174)$ | $(0.498)$ | $(0.207)$ |
| sigmagain | 0.391 | 272.9 | $0.0139^{*}$ |
|  | $(0.298)$ | $(317.0)$ | $(0.00618)$ |
| sigmaloss | 0.486 | 49588.6 | $0.179^{*}$ |
|  | $(0.288)$ | $(73988.2)$ | $(0.0748)$ |
| $N$ | 4738 | 4722 | 4242 |
| AIC | 6017.9 | 5485.8 | 4729.0 |
| ll | -3005.0 | -2738.9 | -2360.5 |

Standard errors in parentheses
${ }^{*} p<0.05,{ }^{* *} p<0.01,{ }^{* * *} p<0.001$
do not affect measured risk preferences for gains or losses - most of the observed significant coefficients are the result of small sample sizes for that demographic group. We do find evidence that women are more risk-averse than men over money gains and leisure losses. However, we find that they are actually more risk-tolerant than men in leisure gains, and no different than men over money losses. Sex differences in risk preferences are have been found in some experiments but not others (see Harrison et al. (2007) for a small discussion), and may be the result of other unobserved characteristics. Chinese-speaking students are more riskaverse over money losses. Strangely, their preferences for leisure losses cannot be effectively described - they appear to be so risk-seeking that it breaks the estimation technique. This may be due to confusion about the task - scores for Chinese-speaking students in the comprehension quizzes were significantly lower than English-speaking students.

TABLE 15. Order Effects

|  | $(1)$ <br> choicesmoney | $(2)$ <br> choicesleisure |
| :--- | :---: | :---: |
| alpha |  |  |
| 2.session | 0.0133 | $-0.304^{*}$ |
|  | $(0.0601)$ | $(0.146)$ |
| 2.framing | 0.0430 | -0.193 |
|  | $(0.0677)$ | $(0.128)$ |
| Constant | 0.339 | $-1.012^{* * *}$ |
|  | $(0.226)$ | $(0.272)$ |
| beta |  |  |
| 2.session | -0.0624 | -0.0702 |
|  | $(0.0640)$ | $(0.112)$ |
| 2.framing | $-0.156^{*}$ | -0.0104 |
|  | $(0.0642)$ | $(0.104)$ |
| Constant | $1.505^{* * *}$ | $2.759^{* * *}$ |
|  | $(0.207)$ | $(0.407)$ |
| sigmagain | 0.409 | $0.0111^{*}$ |
|  | $(0.329)$ | $(0.00536)$ |
| sigmaloss | 8.663 | 219.6 |
|  | $(5.229)$ | $(261.1)$ |
| $N$ | 4738 | 4722 |
| $A I C$ | 6022.2 | 5458.1 |
| ll | -3003.1 | -2721.0 |
| Standard errors in parentheses |  |  |
| $* p<0.05,{ }^{* *} p<0.01,{ }^{* * *} p<0.001$ |  |  |

TABLE 16. Demographic Covariates

|  | (1) <br> choicesmoney <br> alpha |  | choicesleisure <br> blpha |  |
| :--- | :---: | :---: | :---: | :---: |
| Female | $-0.161^{* *}$ | -0.0373 | $0.253^{* *}$ | $0.242^{*}$ |
|  | -0.0574 | -0.056 | -0.0892 | -0.109 |
| Business | 0.0163 | 0.00797 | 0.238 | -0.0191 |
| Journalism | -0.107 | -0.0808 | -0.233 | -0.161 |
|  | -0.05 | -0.203 | 0.376 | -0.172 |
| Social Science | -0.168 | -0.154 | -0.228 | -0.218 |
|  | -0.0474 | 0.0294 | 0.272 | $0.337^{*}$ |
| Natural Science | -0.166 | -0.109 | -0.241 | -0.17 |
|  | 0.0635 | $0.301^{* * *}$ | 0.425 | 0.107 |
| Other | -0.192 | -0.0742 | -0.255 | -0.132 |
|  | 0.0107 | -0.148 | 0.0549 | 0.0191 |
|  | -0.185 | -0.211 | -0.264 | -0.208 |
| Sophomore | -0.0516 | $0.221^{* *}$ | -0.0221 | -0.0323 |
|  | -0.085 | -0.0725 | -0.0971 | -0.132 |
| Junior | 0.0756 | -0.123 | 0.114 | 0.0174 |
| Senior | -0.145 | -0.106 | -0.137 | -0.181 |
|  | -0.0428 | 0.0146 | -0.177 | -0.0204 |
|  | -0.137 | -0.0947 | -0.17 | -0.145 |
| Chinese | 0.151 | $0.262^{* *}$ | $-0.901^{*}$ | -118.4 |
|  | -0.112 | -0.101 | -0.407 | $()$. |
| Other | $0.301^{* * *}$ | -0.175 | -1.461 | -0.445 |
|  | -0.0843 | -0.101 | -2.019 | -0.272 |
| numeracytotal | 0.0613 | $0.107^{*}$ | $0.124^{* *}$ | 0.0575 |
|  | -0.0374 | -0.0487 | -0.0478 | -0.0644 |
| Constant | 0.304 | $0.980^{* *}$ | $-1.742^{* * *}$ | $2.320^{* * *}$ |
|  | -0.339 | -0.303 | -0.381 | -0.428 |
| Sigmagain | 0.511 |  | $0.0101^{*}$ |  |
| Sigmaloss | -0.485 |  | -0.00482 |  |
| Observations | 6.499 |  | 109.1 |  |
| AIC | -3.803 |  | -99.76 |  |
| ll | 4738 |  | 4722 |  |
| Sta | 5986.1 |  | 5184.8 |  |
| -2966.1 |  | -2565.4 |  |  |

Standard errors in parentheses
${ }^{*} p<0.05,{ }^{* *} p<0.01,{ }^{* * *} p<0.001$

From the risk attitudes in different categories that we elicited with the DOSPERT assessment (from Blais and Weber (2006)), we examine any connections between estimated risk preferences in the lab and stated willingness to take various risky actions. Subjects took the 30-question assessment, rating their likelihood on a 7-point scale of taking various risky actions in five categories of risk: ethical, financial, health, recreational, and social. For each subject, the total rating for each category was added up, then normalized over all subjects, so the variables we use are measured in standard deviations. Table 17 shows the estimates using these five covariates. Note that this does not control for demographics - we are investigating the effect of stated risk attitudes, not stated risk attitudes beyond those expected from demographic differences.

We would expect attitudes toward financial risk to be correlated with risk preferences for money, but do not find that - financial risk attitudes have no significant effect. Generally, we find few significant correlations of measured risk preferences with self-reported risk attitudes. Ethical risk tolerance corresponds with risk aversion for leisure, and health risk tolerance with leisure risk tolerance, but these are likely spurious.

Similarly, we investigate the connection between personality: normalized scores from Goldberg's 50-question Big Five assessment (Goldberg (1992)) and estimated risk preferences. We have few strong priors here: while some personality traits seem like they might be associated with risk tolerance or seeking (like openness), the connection is not very clear. Table 18 shows the effects of these covariates. Only two correlations out: increased conscientiousness means subjects are more risk tolerant of leisure gains, and extraversion means risk tolerance for leisure losses. Neither of these align with a priori expectations or narratives, and may again be spurious.

TABLE 17. Risk Attitude Covariates

|  | $(1)$ |  | $(2)$ |  |
| :--- | :---: | :---: | :---: | :---: |
|  | choicesmoney | choicesleisure |  |  |
| ethicalrisk | alpha | beta | alpha | beta |
|  | 0.0485 | 0.00230 | $-0.188^{* *}$ | -0.112 |
|  | $(0.0328)$ | $(0.0262)$ | $(0.0672)$ | $(0.0819)$ |
|  |  |  |  |  |
| healthrisk | $(0.0303$ | -0.0285 | -0.0736 | -0.126 |
|  | -0.0131 | $0.0591)$ | $(0.0506)$ | $(0.0659)$ |
| recreationalrisk | $(0.0392)$ | $(0.0417)$ | $0.167^{*}$ | 0.0402 |
|  | 0.0369 | 0.0368 | -0.0706 | -0.0184 |
| socialrisk | $(0.0374)$ | $(0.0343)$ | $(0.0594)$ | $(0.0579)$ |
|  | 0.0561 | -0.0529 | 0.0782 | 0.0934 |
| Constant | $(0.0353)$ | $(0.0390)$ | $(0.0652)$ | $(0.0555)$ |
|  | 0.462 | $1.353^{* * *}$ | $-1.353^{* * *}$ | $2.582^{* * *}$ |
|  | $(0.245)$ | $(0.237)$ | $(0.306)$ | $(0.372)$ |
| sigmagain | 0.580 |  | 0.00843 |  |
|  | $(0.502)$ |  | $(0.00455)$ |  |
| sigmaloss | 7.631 |  | 151.2 |  |
|  | $(5.454)$ |  | $(162.6)$ |  |
| $N$ | 4738 |  | 4722 |  |
| AIC | 6022.0 |  | 5309.6 |  |
| ll | -2997.0 |  | -2640.8 |  |
| Stand |  |  |  |  |

Standard errors in parentheses
${ }^{*} p<0.05,{ }^{* *} p<0.01,{ }^{* * *} p<0.001$

TABLE 18. Personality Covariates

|  | $(1)$ |  | $(2)$ |  |
| :--- | :---: | :---: | :---: | :---: |
| choicesmoney | choicesleisure |  |  |  |
| agreeableness | alpha | beta | alpha | beta |
|  | $(0.00733$ | -0.0147 | 0.0858 | 0.0463 |
|  |  | $(0.0273)$ | $(0.0561)$ | $(0.0692)$ |
| conscienciousness | -0.0167 | 0.0284 | 0.0546 | $-0.136^{*}$ |
|  | $(0.0415)$ | $(0.0402)$ | $(0.0795)$ | $(0.0664)$ |
| stability | $(0.0530$ | 0.0191 | $0.183^{* *}$ | 0.00592 |
|  | -0.00289 | 0.0379 | 0.0513 | 0.0622 |
| openness | $(0.0495)$ | $(0.0253)$ | $(0.0640)$ | $(0.0723)$ |
|  | 0.0251 | 0.00654 | 0.0832 | 0.154 |
| Constant | $(0.0503)$ | $(0.0403)$ | $(0.0729)$ | $(0.104)$ |
|  | 0.355 | $1.359^{* * *}$ | $-1.073^{* * *}$ | $2.420^{* * *}$ |
|  | $(0.253)$ | $(0.235)$ | $(0.219)$ | $(0.427)$ |
| sigmagain | 0.394 |  | $0.0147^{*}$ |  |
|  | $(0.339)$ |  | $(0.00616)$ |  |
| sigmaloss | 7.749 |  | 100.7 |  |
|  | $(5.490)$ |  | $(116.2)$ |  |
| $N$ | 4738 |  | 4722 |  |
| AIC | 6036.6 |  | 5384.9 |  |
| ll | -3004.3 |  | -2678.4 |  |

Standard errors in parentheses
${ }^{*} p<0.05,{ }^{* *} p<0.01,{ }^{* * *} p<0.001$

TABLE 19. Individual Estimates

|  | mean | sd | count | min | $\max$ |
| :--- | :---: | :---: | :---: | :---: | :---: |
| alpha0money | .502 | 1.501 | 67 | -3.455 | 4.527 |
| alpha0leisure | -.287 | 1.220 | 68 | -3.3179 | 2.944 |
| alpha1money | .387 | 1.436 | 55 | -2.069 | 5.186 |
| beta1money | 1.433 | 1.426 | 55 | -2.068 | 4.016 |
| alpha1leisure | -.506 | 2.001 | 48 | -10.782 | 3.499 |
| beta1leisure | 1.566 | 1.896 | 48 | -3.680 | 5.351 |

### 4.6 Individual Results

Our experiment also allows us to investigate individual risk preferences for both money and leisure, and for gains and losses. Unfortunately, the two-parameter model converges for both money and leisure for only 33 out of 79 subjects. When restricted to the one-parameter model, we have estimates in both money and leisure for 59 subjects. Note that "alpha0" denotes the measured $\alpha$ in the one-parameter model, while "alpha1" and "beta1" are estimates from the two-parameter model. Table 19 shows summary statistics for all the reported risk preference parameters, and Figure 8 shows histograms of all the different measures of risk preference (with a few outliers removed). All of these graphs reflect a wide range of risk attitudes, ranging from very risk-averse to very risk-seeking.

Table 20 shows the correlations between the six estimates of risk preferences, and Figure 6 shows three of these relationships in scatter plots: the one-parameter estimates for money and leisure, and comparing preferences of gains versus losses for both money and leisure. Stars on the table reflect significance of these relationships in a simple linear regression model.

Not surprisingly, we find strong correlations between the one-parameter estimates and the same-domain two-parameter estimates, as these are explaining the same phenomenon. We also find negative correlations between $\alpha$ and $\beta$ for both

FIGURE 8. Individual Histograms


TABLE 20. Individual Correlations

|  | alpha0mon | alpha1mon | beta1mon | alpha0leis | alpha1leis | beta1leis |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: |
| alpha0money | 1 |  |  |  |  |  |
| alpha1money | $0.6493^{* * *}$ | 1 |  |  |  |  |
| beta1money | $-0.7812^{* * *}$ | $-0.1602^{*}$ | 1 |  |  |  |
| alpha0leisure | 0.181 | 0.3407 | -0.0054 | 1 |  |  |
| alpha1leisure | 0.3242 | 0.3622 | $-0.3051^{*}$ | $0.6696^{* * *}$ | 1 |  |
| beta1leisure | 0.0006 | -0.0946 | -0.1219 | $-0.8217^{* * *}$ | $-0.2886^{*}$ | 1 |
| $* p<0.05$ * $^{* *} p<0.01$ **** $^{* *<0.001}$ |  |  |  |  |  |  |

FIGURE 9. Individual Scatter Plots

money and leisure: subjects who are more risk-averse in gains are also more riskaverse in losses. These correlations are significant but relatively small: individuals vary widely in how they view gains versus losses. We also have somewhat large correlations between preferences for money and leisure: individuals who are riskaverse to one seem to be more risk-averse in the other domain. However, these cross-domain correlations are mostly not significant, perhaps due to a small sample size: we only recovered estimates for money and leisure for 33 subjects.

These results show that while there is some connection between individual risk preferences for money and for effort, they are only marginally correlated: most of the individual variation in risk preference for leisure are not predicted by variation in the money domain. Therefore, while risk preferences for money may serve as a weak approximation for risk preferences for effort and leisure, any deeper analysis of behavior over effort requires specific knowledge of preferences for effort.

### 4.7 Conclusion

With this experiment, we have shown that risk preferences for real-effort tasks and risk preferences for money are distinct: they are correlated with different demographics and covariates, and only very weakly predict each other at an individual level. This is particularly surprising: while we find a wide range of risk preferences for both money and leisure in our sample, individuals who are riskaverse over money are only slightly more likely to be risk-averse over leisure. Also, our experiment finds little evidence of the traditional "fourfold pattern" which includes being risk-seeking over losses. Subjects are clearly risk-averse over losses of leisure in all framings, although less so than with gains. Our experimental design does not allow for direct utility comparisons of gains and losses, but we do find differences in utility curvature over gains versus losses for both money and leisure,
even when losses are measured as the final resulting gain. However, we find only a weak connection between risk preferences over gains and over losses, in both the money and leisure domains.

Because we find that risk preferences for money and leisure are only weakly connected and with very different absolute levels, structural assumptions about risk across different domains may not be accurate, and continued research into risk preferences in different domains is warranted. More careful study focusing on what causes these differences across domains may be warranted. Future research may use alternative experimental techniques to further investigate risk preferences for leisure, and possibly investigate risk preferences over combinations of money and leisure - how do subjects view risks across multiple decision-making domains.

## REFERENCES CITED

Abeler, J., Falk, A., Goette, L., and Huffman, D. (2011). Reference points and effort provision. The American Economic Review, 101(2):pp. 470-492.

Andersen, S., Harrison, G. W., Lau, M. I., and Rutström, E. E. (2008). Eliciting risk and time preferences. Econometrica, 76(3):583-618.

Andreoni, J., Kuhn, M. A., and Sprenger, C. (2013). On measuring time preferences. Technical report, National Bureau of Economic Research.

Andreoni, J. and Sprenger, C. (2012a). Estimating time preferences from convex budgets. The American Economic Review, 102(7):3333-3356.

Andreoni, J. and Sprenger, C. (2012b). Risk preferences are not time preferences. The American Economic Review, 102(7):3357-3376.

Appelt, K. C., Hardisty, D. J., and Weber, E. U. (2011). Asymmetric discounting of gains and losses: A query theory account. Journal of Risk and Uncertainty, 43(2):107-126.

Augenblick, N., Niederle, M., and Sprenger, C. (2015). Working over time: Dynamic inconsistence in real effort tasks. The Quarterly Journal of Economics, 1067:1115.

Augenblick, N. and Rabin, M. (Working Paper). An experiment on time preference and misprediction in unpleasant tasks.

Bisin, A. and Hyndman, K. (2014a). Present-bias, procrastination and deadlines in a field experiment. Technical report, National Bureau of Economic Research.

Bisin, A. and Hyndman, K. (2014b). Present-bias, procrastination and deadlines in a field experiment. Technical report, National Bureau of Economic Research.

Blais, A.-R. and Weber, E. U. (2006). A domain-specific risk-taking (dospert) scale for adult populations. Judgment and Decision Making, 1(1).

Bryan, G., Karlan, D., and Nelson, S. (2010). Commitment devices. Annu. Rev. Econ., 2(1):671-698.

Burger, N., Charness, G., and Lynham, J. (2011). Field and online experiments on self-control. Journal of Economic Behavior $\mathcal{E}$ Organization, 77(3):393-404.

Carpenter, J. P., Matthews, P. H., and Schirm, J. (2007). Tournaments and office politics: Evidence from a real effort experiment.

Cubitt, R. P. and Read, D. (2007). Can intertemporal choice experiments elicit time preferences for consumption? Experimental Economics, 10(4):369-389.

Doyle, J. R. (2013). Survey of time preference, delay discounting models. Judgment and Decision Making, 8(2):116-135.

Dreber, A., von Essen, E., and Ranehill, E. (2014). Gender and competition in adolescence: Task matters. Experimental Economics, 17(1):154-172.

Ericson, K. M. M. and Fuster, A. (2011). Expectations as endowments: Evidence on reference-dependent preferences from exchange and valuation experiments. The Quarterly Journal of Economics, 126(4):1879-1907.

Estle, S. J., Green, L., Myerson, J., and Holt, D. D. (2006). Differential effects of amount on temporal and probability discounting of gains and losses. Memory § Cognition, 34(4):914-928.

Estle, S. J., Green, L., Myerson, J., and Holt, D. D. (2007). Discounting of monetary and directly consumable rewards. Psychological Science, 18(1):58-63.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. Experimental economics, 10(2):171-178.

Fischer, C. (2001). Read this paper later: procrastination with time-consistent preferences. Journal of Economic Behavior $\mathcal{E}$ Organization, 46(3):249-269.

Frederick, S. and Loewenstein, G. (2008). Conflicting motives in evaluations of sequences. Journal of Risk and Uncertainty, 37(2-3):221-235.

Frederick, S., Loewenstein, G., and O'donoghue, T. (2002). Time discounting and time preference: A critical review. Journal of economic literature, 40(2):351-401.

Fudenberg, D. and Levine, D. K. (2006). A dual-self model of impulse control. The American Economic Review, pages 1449-1476.

Fudenberg, D. and Levine, D. K. (2011). Risk, delay, and convex self-control costs. American Economic Journal: Microeconomics, 3(3):34-68.

Gill, D. and Prowse, V. (2012). A structural analysis of disappointment aversion in a real effort competition. The American Economic Review, 102(1):469-503.

Goldberg, L. R. (1992). The development of markers for the big-five factor structure. Psychological assessment, 4(1):26.

Gul, F. and Pesendorfer, W. (2001). Temptation and self-control. Econometrica, 69(6):1403-1435.

Harbaugh, W. T., Krause, K., and Vesterlund, L. (2002). Risk attitudes of children and adults: Choices over small and large probability gains and losses. Experimental Economics, 5(1):53-84.

Harbaugh, W. T., Krause, K., and Vesterlund, L. (2010). The fourfold pattern of risk attitudes in choice and pricing tasks*. The Economic Journal, 120(545):595-611.

Harbring, C. and Irlenbusch, B. (2005). Incentives in tournaments with endogenous prize selection. Journal of Institutional and Theoretical Economics (JITE)/Zeitschrift für die gesamte Staatswissenschaft, pages 636-663.

Hardisty, D. J., Thompson, K. F., Krantz, D. H., and Weber, E. U. (2013). How to measure time preferences: An experimental comparison of three methods. Judgment and Decision Making, 8(3):236-249.

Harrison, G. and Rutström, E. (2009). Expected utility theory and prospect theory: one wedding and a decent funeral. Experimental Economics, 12(2):133-158.

Harrison, G. W., Lau, M. I., and Rutström, E. E. Identifying time preferences with experiments: Comment.

Harrison, G. W., Lau, M. I., and Rutström, E. E. (2007). Estimating risk attitudes in denmark: A field experiment*. Scandinavian Journal of Economics, 109(2):341-368.

Harrison, G. W. and Rutström, E. E. (2008). Risk Aversion in the Laboratory, chapter 6, pages 41-196.

Harsanyi, J. C. (1967). Games with incomplete information played by bayesian players, i-iii part i. the basic model. Management science, 14(3):159-182.

Hey, J. D. and Orme, C. (1994). Investigating generalizations of expected utility theory using experimental data. Econometrica: Journal of the Econometric Society, pages 1291-1326.

Holt, C. A. and Laury, S. K. (2002). Risk aversion and incentive effects. American economic review, 92(5):1644-1655.

Holt, D. D., Green, L., Myerson, J., and Estle, S. J. (2008). Preference reversals with losses. Psychonomic bulletin \& review, 15(1):89-95.

Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. Econometrica: Journal of the Econometric Society, pages 263-291.

Kőszegi, B. and Rabin, M. (2009). Reference-dependent consumption plans. The American Economic Review, 99(3):909-936.

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

Koszegi, B. and Rabin, M. (2007). Reference-dependent risk attitudes. The American Economic Review, 97(4):1047-1073.

Kurzban, R., Duckworth, A., Kable, J. W., and Myers, J. (2013). An opportunity cost model of subjective effort and task performance. Behavioral and Brain Sciences, 36(06):661-679.

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

Lay, C. H. (1986). At last, my research article on procrastination. Journal of research in personality, 20(4):474-495.

Luce, R. (1959). Individual choice behavior: a theoretical analysis.
Matthey, A. (2008). Yesterday's expectation of tomorrow determines what you do today: The role of reference-dependent utility from expectations. Jena Economic Research Paper, (2008-003).

McClure, S. M., Ericson, K. M., Laibson, D. I., Loewenstein, G., and Cohen, J. D. (2007). Time discounting for primary rewards. The Journal of Neuroscience, 27(21):5796-5804.

McClure, S. M., Laibson, D. I., Loewenstein, G., and Cohen, J. D. (2004). Separate neural systems value immediate and delayed monetary rewards. Science, 306(5695):503-507.

Mischel, W., Shoda, Y., and Rodriguez, M. I. (1989). Delay of gratification in children. Science, 244(4907):933-938.

Niederle, M. and Vesterlund, L. (2007). Do women shy away from competition? do men compete too much? The Quarterly Journal of Economics, pages 1067-1101.

O'Donoghue, T. and Rabin, M. (2001). Choice and procrastination. Quarterly Journal of Economics, pages 121-160.

Prelec, D. and Loewenstein, G. (1998). The red and the black: Mental accounting of savings and debt. Marketing Science, 17(1):4-28.

Reuben, E., Sapienza, P., and Zingales, L. (2010). Time discounting for primary and monetary rewards. Economics Letters, 106(2):125-127.

Samuelson, W. and Zeckhauser, R. (1988). Status quo bias in decision making. Journal of risk and uncertainty, 1(1):7-59.

Tversky, A. and Kahneman, D. (1991). Loss aversion in riskless choice: A reference-dependent model. The Quarterly Journal of Economics, 106(4):1039-1061.

Van den Bergh, B., Dewitte, S., and Warlop, L. (2008). Bikinis instigate generalized impatience in intertemporal choice. Journal of Consumer Research, 35(1):85-97.

Westbrook, A., Kester, D., and Braver, T. S. (2013). What is the subjective cost of cognitive effort? load, trait, and aging effects revealed by economic preference. PloS one, 8(7):e68210.

