Regulation of Choice Behavior: An Experiment Investigating the Hypothesis That People Bundle Sequences of Expected Rewards

A Dissertation
Presented to Faculty of Commerce at the University of Cape Town
In Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy in Economics

Author: Alexander Schuhr

Supervisors:

1) **Don Ross, Ph.D. –** University of Cape Town
2) **Glenn W. Harrison, Ph.D. –** Georgia State University

Abstract
This thesis discusses reward bundling as a process that enables decision makers to self-regulate their choice behavior. Most empirical work on intertemporal choice has focused on analyzing impulsive choice. Less effort has been dedicated to explanations of how individuals manage to overcome self-defeating behavior. This thesis evaluates the theory of reward bundling. It presents a set of econometric tools that can be employed to investigate whether actual choice behavior is consistent with the theory of bundling. It reports an experiment with human subjects. Reward bundling has been demonstrated in experiments with pigeons and rats. However, no empirical study using salient rewards and sound econometric model estimation has ever been carried out with humans. The present experiment is, therefore, the first that meets the methodological standards of experimental economics and finds evidence consistent with the presence of reward bundling.

To my wife Manouchka and my daughter Samantha

Acknowledgments

This thesis is the result of many years of preparation and scientific work. Its completion was only possible thanks to the effort and the invaluable support of several outstanding contributors.

The research reported in this thesis would not have been possible without the seminal work by Dr. George Ainslie. His groundbreaking contributions opened the way for a scientific approach to the problem of self-control and inspired the experiment reported in this thesis. I also owe Dr. Ainslie thanks for his involvement in designing this research and for his encouraging comments and constructive criticism after reading an earlier draft of this thesis.

I would like to thank Professor Morten Lau for his involvement in planning this research, which led to some important design improvements. I also thank Professor Lau for his advice regarding the econometric analysis of the experimental data.

Many thanks go to Professor J. Todd Swarthout, who wrote the software used in the experiment and supervised the experimental sessions at Georgia State University (GSU). Working with Professor Swarthout gave me invaluable insights into the practice of experimental work and was a great inspiration with respect to the possibilities of software development for economic experiments.

I am deeply grateful to the University of Cape Town for its support of my work and its generous funding of my training at Georgia State University (GSU). Special thanks go to Dr. Nelleke Bak, whose personal effort has facilitated bureaucratic processes many times.

I am greatly indebted to the Center for the Economic Analysis of Risk (CEAR) at GSU. This thesis was only possible thanks to the generous funding of the experiment by CEAR. I have been extremely lucky to benefit from the excellent training and the stimulating intellectual environment that CEAR provided. The staff at CEAR have been immensely supportive in providing assistance and many indispensable resources. In particular, I would like to thank Mr. Thomas Kelly. The experimental design required repeated payoff transfers to subjects over many months and constant availability. Mr. Kelly handled this delicate task with the utmost reliability and professionalism.

The Experimental Economics Center (ExCen) at GSU also contributed significantly. The experimental sessions were carried out in the ExCen laboratory. I am especially grateful to the skilled team of ExCen research assistants. Their reliable help made it possible to run the experimental sessions smoothly and without major problems. Moreover, I owe personal thanks to the ExCen director Professor James Cox for allowing me to attend his graduate class in experimental economics.

I cannot possibly overstate my gratitude to my supervisors Professor Don Ross and Professor Glenn W. Harrison. It would be a flattering thought that the chance to be mentored by such outstanding scholars was somehow based on my own merit. The truth is, however, that they generously offered me a great opportunity that I had yet to earn. When I started my Ph.D. work, I was still far too ignorant to fully comprehend the magnitude of the task ahead and the chance offered to me. Now, years later and presumably somewhat more knowledgeable, I am humbled by the breadth and depth of their expertise and honored by their willingness to pass some of it on to me. For better or worse, a somewhat unusual educational background has exposed me to a large number of professional instructors. I have had the privilege to meet my two finest teachers, in the best sense of the word, at the end of my formal education.

Pursuing a Ph.D. would not have been possible without Professor Ross accepting me as his student, despite the fact that I had no undergraduate background in economics. Although I held a deep fascination with human decision making, I had little serious training in it. Nevertheless, Professor Ross agreed to supervise my studies and introduced me to the economic research on self-control and to the experimental project reported in this thesis. Throughout the years he has patiently supervised my studies, directed my work, and enabled me to collaborate with other outstanding researchers. He has not only given me advice regarding specific projects, but also taught me guidelines to distinguish good from bad science. He has introduced me to a novel understanding of economic agency and to a new outlook on behavioral research. His broad interdisciplinary knowledge and his deep understanding of methodological issues have been a constant source of inspiration. Professor Ross has generously funded much of my training and my opportunity to participate in the work that would result in the present thesis. He has also offered me the opportunity to gain additional funding and practice by working as a research assistant for him. Moreover, he has proven extremely flexible in supervising me, allowing me to simultaneously focus on my research as well as other life-changing developments. I wholeheartedly appreciate all the support he has offered through the years.

Another invaluable pillar of support has been Professor Harrison. Without ever meeting me before, he accepted my involvement in the promising and exciting experiment reported in this thesis. His profound understanding of economic theory and econometric analysis is truly astonishing and he is a humbling yet inspiring example of academic excellence. He has installed in me a sense for the deep interconnectedness of theory, experimental design, and econometric analysis, which will continue to inform all the research work I still hope to do. He has also replaced my untrained, naïve approach to behavioral theories, with a clear understanding that latent constructs must be inferred and estimated and has taught me an indispensable set of tools for this task. Professor Harrison has high standards for scientific work, which he uncompromisingly upholds. Meeting these standards has been challenging at times, but also the best guarantee for the quality of my work and my training that I could possibly hope for. Given my lack of formal training in economics, prior to my Ph.D. studies, teaching me has required much of his goodwill, patience, and pedagogical skills. Professor Harrison has given it all in abundance, inside and outside the classroom: he has taken time out of his busy schedule and even opened his home to me to discuss, review, guide, challenge, and explain. He generously allocated resources of CEAR to realize the experiment reported in this thesis. Moreover, he has been flexible and understanding whenever important personal reasons affected the work schedule. With utmost sincerity, I thank him for all the help and advice he has given.

I owe my deepest gratitude to my wife Manouchka Kelly Labouba. My words of acknowledgement cannot possibly do justice to all the support she has so selflessly granted. She has directly helped me with this thesis by patiently proofreading my various drafts, identifying my countless errors, and suggesting many improved formulations. This painstaking task would have been tedious for anybody, but Manouchka put up with it while being pregnant, taking care of a newborn, and pursuing her own Ph.D. She has spent many hours, which could have been dedicated to her own work, freeing up time for me to make progress and eventually complete this work. However, Manouchka's contributions go far beyond helping me to finish the actual write-up. Completing my Ph.D. work has been an endeavor that, though ultimately immensely rewarding, has also been stressful, at times even frustrating. Meeting the challenges of these past years would not have been possible without her love, support, encouragement, and friendship. She has continuously inspired me and, when it was necessary, forced me to improve myself. Whatever I have accomplished, since the most fortunate day she entered my life, has been made possible by her. I would not be able to enjoy the rewards, could I not share them with her. Now

that the process of completing this work has come to an end, how could I not happily dedicate it to her? This dedication is only an insufficient but sincere declaration of my love, appreciation, and eternal gratitude.

Finally, I want to thank my daughter Samantha Zuri Schuhr. I started the data analysis and the writing process around the time of her birth. Too young to comprehend, she had to sacrifice a lot of quality time with her father, so I could work on this thesis. My hope that she will benefit from the rewards that may result from completing this work has been my strongest motivation. The mere joy of having her in my life has made it so much more pleasant to finalize this task. I gladly dedicate this thesis to her, hopeful that someday she will understand this gesture as an earnest, yet necessarily imperfect, testimony of my love and how proud I am of her.

1. Introduction

Suppose that people reverse their preferences as the availability of the sooner alternative approaches in time. If such choice behavior occurs, we can say that individuals choose *impulsively*. A pattern of impulsive choice may occur when decision makers repeatedly face pairs of smaller but sooner (SS) and larger but later (LL) rewards to choose from. Such individuals are often fully aware that the LL option is preferable as long as the decision and the availability of SS are far away. Nevertheless, as they approach the decision they opt for SS, only to regret their choice afterwards. Many researchers have claimed that such choice behavior is prevalent in various situations. Primary examples for such self-defeating behavior in the field are substance abuse, gambling problems, or other addictions. The same choice pattern seems to be present in behavioral phenomena that are usually less severe, such as procrastination or lapses in people’s diet or exercise regimens. In the laboratory, some have claimed the existence of “anomalies” for the standard economic model, constant or exponential discounting, which supposedly require alternative theoretical approaches to intertemporal choice (Loewenstein & Prelec 1992).

Economists and psychologists have dedicated considerable effort to developing such alternatives. Several models, notably various specifications of hyperbolic discounting, have been proposed to explain impulsivity. Indisputable evidence for these models, however, is rare. Many experimental studies that seemingly explain impulsive choice behavior, e.g. Thaler (1981), base their findings on the use of hypothetical rewards. This practice does not meet the widely accepted methodological standards of experimental economics. Moreover, claims that decision makers are inherently impulsive face the problem of explaining how impulse control is possible. If the inherent preferences of people are such that they result in impulsive choice behavior, one must present a plausible reason why they can usually overcome this impulsivity. Virtually everybody has experienced the urge to succumb to short-term temptations. Nevertheless, most people obviously manage to control such urges. This suggests that choice regulation is at least as puzzling as self-defeating behavior. Self-control does not imply the absence of temptation. Instead, an individual who is said to exercise "willpower" has learned to avoid following his appetite "recklessly." Sometimes people anticipate their urges and take precautions, so that they do not fall prey to their own desires. A classical example, as presented by Ainslie (1974) and Elster (1979), is Ulysses on his way home, who asks his crew to tie him to the mast, in anticipation of his inability to control his passion when he will be tempted by the Sirens. Ainslie (1975, 1992, 2001, 2012) suggests that self-control is usually exercised not by such drastic physical measures, but by a more common psychological operation that we often apply when we restrain ourselves in more mundane situations: *reward bundling*.

Ainslie proposes an increase in the tendency to choose the LL reward over the SS reward if individuals face a series of SS/LL pairs, rather than a single SS/LL pair in isolation. He argues that people may *bundle* the rewards in the series, by regarding their current choice as a test case for future choices. It is the aggregation of rewards that helps them overcome the impulse to choose SS. Ainslie (1974, 1975) and Ainslie and Monterosso (2003) find choice patterns displayed by pigeons and rats that are consistent with the reward bundling hypothesis. However, experimental studies of reward bundling with human subjects are rare. The few available examples do not meet the methodological requirements of experimental economics. Hypothetical rewards are not sufficient to establish credible evidence. Moreover, even if real rewards are used, many economists feel uncomfortable with instructions, such as "the choice you make now is the best indication of how you will choose every time" (Kirby and Guastello 2001, 159). Such phrasing could be potentially problematic, given another widely accepted guideline for economic experiments: the nonacceptance of subject deception.[1] Given these considerations, the experiment presented in this thesis is the first to meet the methodological standards of experimental economics.

In chapter 2 several models of time preferences are presented, which are considered in the subsequent econometric analysis of the experimental data. The initial focus is on exponential and hyperbolic discounting, since they give a good overview of the main competing views of discounting behavior that have been proposed. Then additional models are discussed. Chapter 3 discusses some phenomena that are frequently raised in the literature, namely non-linear utility, declining discount rates, and the "magnitude effect," that motivated the experimental design and the econometric analysis. Some considerations from the literature in economics and cognitive psychology are presented. They show how mental representations and "rounding behavior" may result in seemingly declining discount rates and an apparent magnitude effect. A presentation of the theory of reward bundling is presented in chapter 4. The thesis then presents a discussion of the econometric approach to the analysis of the experimental data. Chapter 5 contains a brief discussion of the descriptive methods used: local polynomial regressions and descriptive probit estimations. The emphasis, however, is on the foundations of the structural econometric analysis, presented in chapter 6. Full information maximum likelihood estimation is introduced as a method to obtain estimates of the latent processes that generate the observable choice behavior and to test whether it is consistent with the reward bundling hypothesis. The thesis proceeds to

[1] There is no intended implication that Kirby and Guastello (2001) seek to deceive their subjects. This is not common-or-garden deception of the kind that an economist critic might usually condemn - it is not *deliberate* deception. But it reports to subjects as a fact something that is not established, and indeed likely is not a fact. Hence, the phrasing of their "suggested condition" is problematic from the perspective of experimental standards in economics.

present the experiment which tests for the presence of reward bundling. Chapter 7 lays out the experimental design and procedures. Chapter 8 reports the results and chapter 9 concludes the thesis.

2. Time preferences and reward bundling

2.1 The costs of delay and temporal weighting

People generally prefer receiving a good now over receiving the same good at a later time, all else equal. This tendency is not inherently problematic. In fact, economic theory can well account for this observation and predicts that individuals devalue future goods in a consistent fashion. We can think of discount factor D as a theoretical construct that equates the present value of a reward with its future value for a given time horizon τ. The concept of value employed here is the economists' notion of *utility*. Here is a general model that relates two subjective values of a reward x received at two different points in time:

$$U(x_t) = DU(x_{t+\tau}), \qquad (0)$$

where $U(x_t)$ is the utility of the reward x, if received at time t, and $U(x_{t+\tau})$ is the utility of the reward x, if received at the later time $t + \tau$. The model is general in the sense that it does not assume any specific functional form, either for the instantaneous utility function U or the discount function D. Most models of time preference specify the discount factor D in a way that imposes costs for delaying the reward x from t to $t + \tau$. The rate of devaluation per time unit is determined by a discount rate d. Depending on how the costs are specified, the implied choice behavior is either temporally consistent or involves preference reversals, for reasons to be made clear below. Two archetypical model types that describe either choice behavior are *exponential* and *hyperbolic* discounting.

The *exponential discounting* model has become the standard way of approaching the economics of time preferences. It was introduced by Samuelson (1937) as the *discounted utility model*, although he did not propose it as an empirical hypothesis. Samuelson (1937, 156) describes the essential characteristics of the model as follows: “(t)he individual discounts future utilities in some simple regular fashion which is known to us. For simplicity, we assume in the first instance that the rate of discount of future utilities is a constant.”

The apparent occurrence of “anomalies” for exponential discounting has led many researchers to search for alternatives. The original model to explain apparent preference reversals and self-defeating behavior is the hyperbolic discounting model. Ainslie (1975) initially proposed a simple functional form to describe hyperbolic discounting, but several other specifications have been suggested. Subections 2.2.1 and 2.2.2 discuss exponential and hyperbolic discounting in some detail.

Three additional models are considered in subsection 2.2.3. The *quasi-hyperbolic discounting* model has become popular in economics (Phelps & Pollack 1968, Laibson 1997). Initially introduced for its analytical convenience, it has become associated with the notion that preference reversals result from the *visceral* temptation of immediately available rewards. Quasi-hyperbolic discounting assumes that the costs of delay come in form of a fixed proportion of the principal, as well as a constant cost of delay that varies with the horizon as in the exponential discounting model. Finally, a model is considered that does not approach delay as imposing costs. Instead, it assumes that individuals perform *temporal weighting*, in the sense that their subjective perception of time is distorted. All models presented in this chapter are discussed in detail by Andersen, Harrison, Lau, and Rutström (2011).

2.1.1 Exponential discounting

Samuelson's (1937) model followed a line of research that relied heavily on the discussion of the psychological motives behind intertemporal choice. This literature started as early as 1834 with the publication of John Rae's *The Sociological Theory of Capital* (Rae 1834). Rae employed a psychological variable, the "effective desire of accumulation," to explain why the propensity to delay immediate consumption differs among nations. This psychological reasoning was continued by others. For example, William Jevons (1888) and his son Herbert Jevons (1905) insisted that only the anticipation of greater consumption in the future can help overcome the urge for immediate consumption. Eugen von Böhm-Bawerk (1889) contributed another psychological explanation – one that principally allows for the possibility of self-defeating behavior. He introduced a technical approach that treated intertemporal choice as the self-allocation of resources at different points in time. Irving Fisher (1930) further formalized this approach. He too continued the deliberation of possible psychological motives that may underlie the decision to delay rewards that are immediately available. Historical developments leading to the exponential model are reviewed by Frederick, Loewenstein, O'Donoghue (2002, 352-355).

Unlike his predecessors, Samuelson (1937) did not engage in detailed discussions about the psychology of intertemporal choice. All possible motives, which may affect if and to what extent future rewards are devalued, are compressed in one single parameter δ, the discount rate. In the general model in equation (0) it was shown that the discount factor D equalizes the utilities of a reward received at time t and the same reward received at a later time $t + \tau$. In the exponential model the discount factor is

$$D^E(\tau) = 1/(1+\delta)^\tau, \tag{1a}$$

in discrete time units. The specification in continuous time illustrates why the model is called "exponential"

$$D^E(\tau) = \exp(-\delta\tau). \tag{1b}$$

Like most discounting models, the exponential model states that time delay is costly, i.e. the utility of a reward received later is diminished compared to the same reward received sooner. Unlike other models, however, exponential discounting assumes that the decline in utility occurs at a constant rate $d = \delta$. For example, the percentage rate by which utility is diminished between today and tomorrow is the same as the rate by which utility is discounted between one year and one year and one day from today. One can think of the discount rate as representing the variable costs of time delay. Accordingly, exponential discounting proposes that time delay does not impose any fixed costs, which would be independent of the duration of the delay. Instead, the costs of delaying a reward vary constantly with the time horizon. As a result, it is normal for an economist to view the exponential discounting model as using a constant variable utility cost of time delay and a zero fixed utility cost of time delay.

Strotz (1956) demonstrated that only exponential discounting prevents inconsistent discounting over time. Other models, such as hyperbolic discounting discussed in the following subsection, raise the possibility of continuous re-evaluations and changes of choices.

2.1.2 Hyperbolic discounting

The apparent evidence for repeated preference reversals motivated researchers to find alternatives to the exponential model. The research that would eventually lead to the hyperbolic discounting model began with psychological experiments on the relative allocation of behavior to reward alternatives. Early experiments were performed with animal subjects. Herrnstein

(1961) conducted an experiment in which pigeons chose between two possible reward schedules by pecking keys in a Skinner box. He found that their relative frequency of choice approximately equaled the relative frequency of reinforcement of each schedule. In the psychological literature, Herrnstein's finding is known as the *matching law*, since the relative rate of responding on a key *matches* the relative rate of reinforcement. Inspired by the original matching law, additional factors that affect the relative distribution of behavior were examined. Chung and Herrnstein (1967) studied how delay affects choices and found that the relative frequency of responses matches the relative immediacy of reinforcement. The introduction of delayed rewards drew attention to another issue: subjects prefer larger rewards over smaller ones and sooner rewards over later ones. Obviously, these two tendencies are in conflict when subjects have to choose between smaller, sooner (SS) rewards and larger, later (LL) rewards. Ainslie pointed out that the matching law implies reversal of preference between LL and SS rewards as a function of delay, which might create an incentive for early precommitment to LL alternatives. This precommitment was observed using both differential rates of pecking (Rachlin & Green, 1972) and single choices (Ainslie, 1974).

Ainslie (1975) proposed that time preferences are best described by the *hyperbolic discounting model*. Hyperbolic discounting still assumes variable costs of delay. However, unlike exponential discounting, hyperbolic discounting does not assume that those variable costs are constant. Ainslie proposed a specification of the discount factor D that describes this non-constant pattern

$$D^{H1}(\tau) = 1/\tau \qquad (2)$$

and implies the discount rate

$$d^{H1}(\tau) = \tau^{(1/\tau)} - 1. \quad (3)$$

Other specifications of the hyperbolic model have been suggested. A commonly used specification is due to Mazur (1987):

$$D^{H2}(\tau) = \frac{1}{1+K\tau}. \quad (4)$$

for some parameter K, implying the discount rate

$$d^{H2}(\tau) = (1 + K\tau)^{(1/\tau)} - 1. \quad (5)$$

An individual whose time preferences are described by hyperbolic discounting will show declining discount rates with an increasing horizon. That is, the closer the reward is in time, the more impatient will such an individual be. More importantly, such preferences can account for preference reversals, since an individual may re-evaluate the preferences order of two rewards and change his choices as time passes by.

2.1.3 Other models

Another deviation from the exponential model is *quasi-hyperbolic discounting*. The model was originally introduced by Phelps and Pollack (1968), to model intergenerational transfers. It was rediscovered and popularized by Laibson (1997). The quasi-hyperbolic model, in discrete time units, is given by

$$D^{QH}(\tau) = 1, \quad (6a)$$

at $\tau = 0$ and

$$D^{QH}(\tau) = \beta/(1+\delta)^{\tau}, \quad (6b)$$

for $\tau > 0$. The quasi-hyperbolic model has a jump-discontinuity in the discount factor at $\tau = 0$ and behaves thereafter like the exponential model. The model implies the discount rate

$$d^{QH}(\tau) = [\beta / (1 + \delta)^{\tau}]^{(-1/\tau)} - 1. \qquad (7)$$

The parameter β introduces a fixed cost component that is not present in exponential discounting. It takes a value of $0 < \beta < 1$, which reduces the utility of the reward by a fixed proportion of the principal, as soon as any delay occurs. After this initial effect wears off, the discount rate converges towards δ.

The initial reason for applying the quasi-hyperbolic model to individual time preferences was that it “mimics the qualitative property of the hyperbolic discount function, while maintaining most of the analytical tractability of the exponential discount function” (Laibson 1997, 450). Hyperbolic specifications display declining discount rates with an increasing horizon. Quasi-hyperbolic discounting seeks to mimic this qualitative property with the initial drop in value, due to the parameter β. This drop can be interpreted as representing the costs of delay. That is, as soon as a reward is delayed its value decreases by a proportion of the principal.[2] Unlike hyperbolic specifications the quasi-hyperbolic model maintains the analytical tractability of exponential discounting, since it preserves the use of limit theorems to predict unique preferences from knowing an individual’s prior incentives (Ainslie 2012, 8). Since Laibson (1997) reintroduced the quasi-hyperbolic functional form, however, it has undergone a reinterpretation, which will be discussed in section 3.2.

The final discounting model considered in this chapter deviates from the others in the sense that it does not propose fixed or variable costs of delay. Instead, it regards individuals as

[2] These costs do not necessarily have to come in the form of a fixed percentage of the principal. Benhabib, Bisin, and Schotter (2010) propose a fixed costs model that introduces present bias by assuming fixed monetary costs of delay. In their model, delay is costly because it decreases the value of the reward by a fixed monetary amount. For example, an individual may not be willing to accept delay, unless he is compensated by $1, or $5, or another specific amount. This model is pedagogically valuable, because it demonstrates that the popular quasi-hyperbolic model is only one possible way to describe present bias. However, Andersen et. al. (2014b) find no evidence for fixed monetary costs of delay.

perceiving time in a distorted way. This is comparable to probability weighting in decision making under risk. In an intertemporal context, people are assumed to "speed up" or "slow down" time in their perception. A flexible specification that describes this behavior is due to Read (2001, 25), based on the Weibull distribution:

$$D^W(\tau) = \exp(-\delta\tau^{(1/s)}), \tag{8}$$

for $\delta > 0$ and $s > 0$. The parameter s is responsible for the "speeding up" or "slowing down" of perceived time. The specification collapses to the exponential model if $s = 1$.

3. Motivation for design and analysis of the control treatments

Theory, experimental design, and econometric analysis are deeply connected. Consequently, the treatments and analytical strategies presented in the subsequent chapters are informed by theoretical considerations. The treatment conditions, implemented in the present experiment, are presented in Table 1. This chapter discusses the motivation for four control treatments that are implemented in the present experiment.

Table 1: Experimental design

Treatment	Description	Date of the session	Number of subjects
T0	The choice between a smaller amount "today" or a larger amount in 28 days.	8/27/2013	35
T1	The same as T0, but with a front-end delay of 1 day.	8/28/2013	29
T2	The same as T0, except the amounts are tripled.	8/29/2013	36
T3	The same as T0, but with a front-end delay of 35 days.	8/28/2013	35
T4	Two "free" decisions. One as in T0, one as in T3.	8/30/2013	37
T5	Two "forced" decisions. One as in T0, one as in T3.	9/3/2013	33
T6	Three "free" decisions. One as in T0, one as in T3, and another pair of options.	9/4/2013	31
T7	Three "forced" decisions. One as in T0, one as in T3, and another pair of options.	9/5/2013	30

Behavioral economists have claimed the existence of several empirical phenomena in intertemporal choice, many of which supposedly constitute anomalies for exponential discounting. The discussion will focus on three claims that are frequently found in the literature. Many previous studies have found that their subjects' choices indicate extremely *high levels of discounting rates*. Section 3.1 argues, following Andersen, Harrison, Lau, and Rutström (2008), that these high discount rates result from neglecting subjects' utility function and specifically the evidence for diminishing marginal utility. It is explained that if the econometric analysis of choice data accounts for the concavity of utility functions the inferred discount rates are lowered drastically.

Two other frequent findings from behavioral economics are that *discount rates decline* with *time* as well as with the *magnitude* of the rewards. Both effects have been prominently discussion in the literature as anomalies of the exponential discounting model (Loewenstein & Prelec 1992, 574-575). These phenomena would affect the probability of LL choices, which is why their possible prevalence is important for the main objective of the experiment: testing whether reward bundling is employed as a means of impulse control. Three control treatments T0, T1 and T3 introduce varying front end delays (FED) to test whether elicited discount rates decline as present decisions are made about rewards that are paid out at different times in the future. Section 3.2 further explains the motivation for these treatments. An additional treatment T2 varies the magnitude of the reward to examine the effect on the implied interest rates. Section 3.3 briefly summarizes the previous literature on the magnitude effect and further explains the motivation for treatment T2.

Section 3.4 discusses how the similarity of SS and LL rewards may lead to rounding behavior that can lead to observations of all three phenomena discussed in the previous section:

very high discount rates, their decline with time, and an apparent magnitude effect. The psychological notion of *similarity* as the basis of rounding is explained. The design implications are that the stakes in the decision task must include amounts of money that are sufficiently high, so that the quantities that are being compared are not too similar and invite rounding.

3.1 Non-linear utility

Risk preferences and time preferences are two fundamental elements of individual decision making. However, it can be difficult to distinguish between their respective effects, since "the future is inherently risky." This can be problematic if one wishes to understand the separate impact of time preferences on decision making (Andreoni, & Sprenger 2012, 3357). It is possible that what appears to be a preference for sooner rewards may, instead, be an aversion to the risk that is inherent in future payoffs.

Risk aversion is described by concave curvature of instantaneous utility functions. Some doubt has been raised whether concave utility functions are plausible, over the stakes that are typically used in laboratory experiments. Following Hansson (1988), Rabin's (2000) calibration theorem suggests that when subjects exhibit risk aversion over "small stake" lottery choices in the laboratory, it is not plausibly explained by concave utility functions. Rabin (2000) demonstrates that the rejection of the small scale lotteries over a wide range of wealth levels entails absurd levels of risk aversion at some wealth level. He argues that explaining small scale risk aversion with concave utility functions implies that the subjects' marginal utility in wealth is declining so rapidly that even lotteries with huge possible gains are implausibly rejected. Rabin and Thaler (2001) interpret this as evidence against expected utility theory (EUT) and argue that an alternative theory is needed to explain risk aversion over small stakes. They base their argument on the assumption that the integration of income into "lifetime wealth" is an essential characteristic of EUT. Cox and Sadiraj (2006) point out that this only applies to one specific version of EUT – the "EUT of terminal wealth," where current income is completely integrated in lifetime wealth. However, nothing in the axiomatic structure of EUT dictates that the argument of utility functions is necessarily terminal wealth. At the opposite end of the theoretical spectrum, current income is valued in isolation and not at all integrated in lifetime wealth, as originally proposed by Vickrey (1945). Cox and Sandiraj (2006) call this alternative model the "EUT of income" (Cox & Sadiraj 2006, 47-52). To what extent people exhibit asset integration is ultimately an empirical question.

Andersen, Cox, Harrison, Lau, Rutström and Sadiraj (2011) distinguish between EUT models with full asset integration (FAI), no asset integration (NAI), and partial asset integration (PAI), with FAI and NAI as extreme cases. They show that only FAI models are affected by Rabin's (2000) critique. They make use of a unique data policy in Denmark, which allows matching personal wealth data with choice data from a risk preference task. The latter were collected in an *artefactual field experiment*, according to the terminology proposed by Harrison and List (2004, 1014). Andersen et.al. (2011) find no support for the terminal wealth model. However, there is evidence for partial asset integration. A fraction of the experimental gains is integrated into total wealth. This supports the PAI model and suggests that subjects can show low-level risk aversion, while simultaneously not being absurdly risk averse in the large. This shows that risk aversion in experimental choice data can be plausibly described by concave utility functions. This finding does not only apply to EUT, but also to alternative models, such as rank-dependent utility (RDU) or cumulative prospect theory (CPT), which attribute some of the observed risk preferences to concave utility.

Consider the general form of a discounted utility model, as presented in (0). If the utility function $U(\cdot)$ is concave in x, it follows from Jensen's inequality that the implied discount rate, for given observed choices over SS and LL options, decreases. This formalizes the intuition that both diminishing marginal utility and discount rates lower the influence of future payoffs on present utility. This has important implications for the elicitation and estimation of individual discount rates. Andersen, Harrison, Lau, and Rutström (2008) show that the commonly performed isolated elicitation and estimation of discount rates is problematic. They elicit risk and time preferences from the same subjects in their artefactual field experiment in Denmark. They combine data from a risk aversion and a time preference task, before performing joint estimation of utility functions and discount rates, using full-information maximum likelihood (FIML).

Previous studies on intertemporal choice usually assume linear utility functions, with the result that inferred discount rates were often extremely high.[3] As one allows for concave utility functions, the implied discount rates decrease. Andersen, Harrison, Lau, and Rutström (2008, 600-603) find annual discount rates of 10.1 percent. This value is implied by an EUT specification. However, the main argument is more general: all concave utility specifications, not only those implied by EUT, result in lower implied discount rates. Moreover, discount rates cannot be estimated based on time preference experiments alone. Ideally, one should conduct separate tasks to elicit the utility functions.[4]

If, for budgetary or efficiency reasons, a risk aversion task cannot be performed in the same experiment, an alternative strategy is available. One can then elicit time preferences

[3] A compact overview of the literature is given in Andersen, Harrison, Lau, and Rutström (2014b, Table 6).
[4] In principle, one can attempt to elicit risk and time preferences in one single task. Recent attempts to accomplish this have been proposed by Andreoni & Sprenger (2012) and Laury, McInnes, & Swarthout (2012). If successful, such a combined elicitation may have desirable budgetary and implementation advantages. However, there are disadvantages that are perhaps less obvious. For example, such a combined elicitation will almost certainly result in a task that is cognitively more demanding for subjects than separate elicitation. The desire to avoid "noisy" estimates of both risk and time preferences is, therefore, and additional reason to conduct separate tasks.

separately provided that estimates of risk attitudes from the same population are available. In this case, inferences about time preferences can still be appropriately conditioned on those risk estimates. The latter strategy is used in the experiment reported in this chapter. Various experiments on risk attitudes, with Georgia State University (GSU) students as subjects, had been performed before the present experiment took place. These data are readily available and are combined with the elicited data on time preferences. Hence, joint estimation of risk and time preferences is possible, even though the budget of the experiment could be completely allocated to the various discounting treatments.

3.2 Declining discount rates

The previous section discussed the need to control for a-temporal *risk preferences*, in order to infer utilities correctly, and implied design requirements. This section focuses on *time preferences*. Although one must control for a-temporal risk attitudes to correctly infer discount rates, *time preferences* and *risk preference*s are different concepts. Nevertheless, both concepts are related in the following sense: When intertemporal choice is over time-dated lotteries, rather than certain payouts, people often display *correlation aversion*, which is also called intertemporal risk aversion by some. This concept and its implications are discussed by Andersen, Harrison, Lau, and Rutström (2011). The following example illustrates the notion of intertemporal utility: suppose subjects face two lotteries, with two time-dated payouts, where the first payout is realized with probabilities p and the second payout is realized with probability $1 - p$. The payouts involve a set of sooner payouts $\{x, X\}$, where $x < X$, and a set of later payouts $\{y, Y\}$, where $y < Y$. Suppose, subjects face a choice between $L_1 = \{0.5[px, (1-p)Y]; 0.5\ [pX, (1-p)y]\}$ and $L_2 = \{0.5[pX, (1-p)Y]; 0.5\ [px, (1-p)y]\}$. So lottery L_1 offers a 50 percent chance of the "bad" outcome sooner and the "good" outcome later and a 50 percent chance of a "good" outcome sooner and a "bad" outcome later. Lottery L_2 offers a 50 percent chance of the "good" outcome sooner and the "good" outcome later and the 50 percent chance of a "bad" outcome at both times. Only when subjects are indifferent between $L_1$and L_2 is the intertemporal utility characterized by an *additive* specification. An additive intertemporal utility function $U^{it}(x, y) = D^S U(x) + D^L U(y)$ defines intertemporal utility as the weighted sum of the discounted instantaneous utility of the sooner and later outcomes. The components x and y are typical sooner and later outcomes, and $D^S$and $D^L$are their respective discount factors. When subjects' choices imply $L_1 \succcurlyeq L_2$ they display correlation aversion, and when $L_1 \preccurlyeq L_2$, then they are correlation loving. Such intertemporal preferences are not correctly characterized by an additive intertemporal utility function. However, a simple extension of the intertemporal model, such as $U^{it}(x, y) = \frac{[D^S U(x)+D^L(y)]^{(1-\xi)}}{(1-\xi)}$, introduces intertemporal risk attitudes and accounts for non-additivity. It is important to be aware that intertemporal correlation aversion does not imply instantaneous risk aversion or *vice versa*. A correlation averse decision maker may be a-temporally risk averse, risk loving, or risk neutral. Intertemporal and instaneous risk aversion are distinct concept, although they are theoretically related.[5]

Various discounting models have been proposed to describe time preferences. Some specifications, notably hyperbolic and quasi-hyperbolic discounting, predict that discount rates decline with the time horizon. Several treatments were introduced to test this prediction. Varying the front-end delay (FED) – the delay until the sooner reward will be paid out – allows testing some of the discounting models presented in section 2.1. Different discounting models imply

[5] This theoretical connection has important econometric implications. Just as controlling for risk attitudes is necessary to correctly infer discount rates, correlation attitudes must be estimated conditional on risk *and* time preferences. This requires full-information maximum likelihood estimation.

different discount rates, as SS rewards are delayed. This is why the experimental design implemented treatments with varying FED.

The basic treatment condition T0 did not use a FED and the subjects received the selected payouts on the day of the experimental session. This procedure came with the potential problem that subjects could perceive immediate payouts as different from future ones. In this case, the elicited time preferences may be confounded by other factors that affect choice behavior.

A practical problem with time-dated rewards is that subjects may perceive immediate rewards as more credible than future rewards. If there are doubts that future rewards will actually be paid out, choosing the sooner option might reflect this uncertainty, rather than time references. To mitigate this issue subjects in our experiments were given a written confirmation, signed by the principal investigator, which guaranteed that they would receive future payment on a specified date. They were also provided with contact information that would enable them to direct questions regarding payments to the experimenters.

Another related problem is that the transaction costs between immediate and later rewards may differ. Immediate rewards can be paid out effortlessly as soon as the experimental session is complete. Obtaining future rewards may come at greater costs. The design employed in the present experiment addresses these concerns by using the same payment method for immediate payouts as for future payouts. Immediate payouts as well as future payouts were transferred using the online payment service *PayPal*. However, immediate payouts were sent later on the same day the experimental session took place (see detailed explanations below).

When transaction costs, credibility issues, and general uncertainty are accounted for, discount rates varying with FED may still indicate deviations from exponential discounting, such as hyperbolic or quasi-hyperbolic time preferences. Coller and Williams (1999) estimate the effect of a one month FED and find that it reduced the elicited discount rates drastically. To test the effect under comparable conditions, a 35 day FED was implemented in treatment T3. Subjects were informed that if they chose the sooner option and the payout is selected, the payout would be sent 35 days after the session via *PayPal*.

If declining discount rates are observed, this may be regarded as evidence for both hyperbolic ($D^{H1}$and D^{H2}) and quasi-hyperbolic (D^{QH}) specifications. However, the difference between these alternative models is rather important. There is a school in behavioral economics that insists that the quasi-hyperbolic model is more than just a convenient approximation of hyperbolic discounting. Researchers from this tradition assert that the quasi-hyperbolic specification reflects actual cognitive processes in intertemporal choice. This position claims that more deliberate reasoning, represented by δ, is often overwhelmed by visceral factors, represented by β. Accordingly, the viscerality of an imminent reward causes the impatience or "passion for the present" (Loewenstein 1996). This "present bias" takes the form of a decline in value for any delay. Such fixed costs of delay can be described by a fixed absolute decline in value, as proposed by Benhabib, Bisin, and Schotter (2010), or as a fixed proportional decline, as described by quasi-hyperbolic discounting. Following the quasi-hyperbolic school of thought, McClure et.al. (2004) even attempted to locate visceral β-discounting and deliberate δ-discounting in different discounting types in different brain areas, using functional Magnetic Resonance Imaging (fMRI).

However, these results could not be replicated by another fMRI study, conducted by Glimcher et.al. (2007). The neuroeconomic results are not the only claims associated with the D^{QH}model that are fiercely contested. The relevance of viscerality in time preferences and the validity of the β parameter have been called into question. Andersen, Harrison, Lau, and

Rutström (2014b, 26-30) do not find support for the D^{QH} model in the adult Danish population. Ainslie (2012, 10-11) discusses empirical phenomena which seem to result from declining discount rates and preference reversal, but are difficult to explain by visceral arousal. A primary example is procrastination, a behavioral pattern that can be well explained by preference reversal but does not involve visceral arousal. The D^{QH} model is also the basis for several models of impulse control, which assume that what must be overcome is "visceral arousal" and "passion for the present." If the D^{QH} model is not supported, these models lack theoretical plausibility (Ainslie 2012, 21-28).

The quasi-hyperbolic model makes two distinct predictions that offer an opportunity to decide whether it describes discounting accurately, or whether the observed choices are better described by hyperbolic discounting:

- The D^{QH} model predicts an immediate decline in utility by a fixed proportion, as soon as there is a delay of the payout. If this decline is interpreted as the effect of visceral arousal it is only plausible to assume its effect can only last as long as such arousal can be plausibly maintained. Suppose that the utility of rewards becoming available soon after this period (e.g. one day after the decision is made) does not differ significantly from the diminished utility that could be explained by visceral arousal. In this case, choice behavior is better explained by hyperbolic discounting.
- the quasi-hyperbolic discounting model D^{QH} also predicts that after the initial impatience wears off, the model approaches exponential discounting. Suppose the discount rate is found to be variable, as a function of delay, even after an extended FED of several days. This observation would also be better explained by hyperbolic discounting.

Two experimental conditions, which were initially introduced to serve as control treatments when the presence of reward bundling is investigated, also enable the analysis to investigate the underlying discounting behavior. Treatment T1 adds a FED of 1 day. If discounting behavior in treatment T1 did not differ significantly from discounting in treatment T0, this would support hyperbolic discounting instead of the visceral interpretation of the quasi-hyperbolic specification. If visceral arousal was an important factor, the elicited discount rate should be significantly different to the discount rate from treatment T0.[6]

Moreover, suppose choices in treatment T3 were more patient than choices in T1. This would be the case if there were significantly more LL choices in T3 than in treatment T1. Such choice behavior would lend additional support for hyperbolic discounting and provide evidence against the proposed role of visceral arousal.

3.3 The magnitude effect

The *magnitude effect* is the observation that subjects tend to display lower discount rates when payoffs are larger. A significant magnitude effect would be an anomaly for exponential discounting and would indicate that discounting functions, and their implied discount rates, cannot be used independent of the scale of the quantities involved (Andersen, Harrison, Lau, and Rutström, 2013).

Many behavioral economists consider the magnitude effect a well established result. Loewenstein and Prelec (1992, 557) claim that "[e]mpirical studies of time preferences have also

---

[6] McClure et.al. (2007, 5796-5797) suggest a modification of the D^{QH}specification introducing an additional horizon parameter τ that determines after what delay the jump discontinuity occurs and that differs for each individual. One could argue that the analysis in the present experiment performs pooled estimation and that such an individual parameter remains undetected. However, even the existence of such an individually varying parameter τ could not invalidate evidence against the D^{QH}in the present design. The reason is that the 1 day FED in treatment T1 lasts longer than visceral arousal can plausibly be maintained.

found that large dollar amounts suffer less proportional discounting than do small ones." Scholten and Read (2010, 929) even describe the magnitude effect as "probably the most robust anomaly in intertemporal choice." Many of the studies that supposedly establish the magnitude effect use hypothetical payoffs. Thaler (1981, 205), for example, uses hypothetical payoffs and concludes that "[a]s the size of the reward increases, the implicit discount rate falls." He also provides a possible explanation for the magnitude effect, suggesting that the mental effort that the postponement of a reward requires is only worthwhile for rewards that are high enough. In other words, the stakes must be worth the wait. The use of hypothetical payoffs is commonly not considered sufficient to meet the requirements of experimental economics. However, Thaler's (1981) results are widely cited and have influenced numerous studies in behavioral economics. A detailed literature review is offered by Andersen, Harrison, Lau, and Rutström (2013, 671-681).

Andersen et.al. (2013) perform a test for the presence of a magnitude effect in an artefactual field experiment with real payouts. The subjects are a representative sample of the adult Danish population. The estimation of discount rates assumes the D^E specification and accounts for non-linear utility functions, as explained in subsection 2.1.1. The exponential discounting function is used in a descriptive, rather than a structural way. The main objective is to test the effect that varying principals have on the implied discount rates. If higher principals generated lower implied discount rates than lower principals, it would suggest the presence of a magnitude effect. Andersen, Harrison, Lau, and Rutström (2013, 20-21) find only a minuscule magnitude effect: the difference between high and low stakes is only a difference of 0.3 percentage points, significant with a *p*-value of 0.004. That is, the slight magnitude effect they find is statistically significant, but economically insignificant for most conceivable policy purposes.

These results indicate that the magnitude effect is not as robust a phenomenon as previous studies suggest. The experimental design discussed in this thesis introduced a control treatment T2 that tripled the principal compared to the baseline treatment T0. The main purpose of this treatment is to control for the fact that *bundling treatments*, which are described below, effectively yield payouts that are two or three times the amount of the control treatments T0, T1, or T3. However, the introduction of treatment T2 also enables the analysis to test for a magnitude effect. If the elicited discount rates from treatment T2 were lower than the discount rates from treatment T0, this would suggest the presence of a magnitude effect.

3.4 Similarity and Rounding

Declining discount rates and the magnitude effect, as found by previous studies, can be produced by subjects' attempt to handle the cognitive burden of evaluating similar quantities. Decision makers often perform mental operations seeking to simplify choice problems, in order to reduce the deliberation cost of optimizing. One possible operation of this kind is rounding behavior, which may result in the apparent occurrence of anomalies from the perspective of exponential discounting. This can be illustrated by a simple example, following Andersen, Harrison, Lau, and Rutström (2013, 682-683). Suppose discount rates are elicited using an open ended "fill-in-the-blank procedure," as employed by Benhabib, Bisin, and Schotter (2010). This method requires that subjects state the amount of money paid in the future that would make them indifferent to a given amount of money paid today. Column (1) of Table 2 includes principals ranging from $5 to $30 in $5 increments. Assume that a representative subject has a discount rate of 10%, which is listed as the annualized effective rate (AER) in column (4). The varying horizons, four multiples of a 28 days horizon, are listed in column (3). The LL amounts that would make a subject, with the given discount rate, indifferent to the principal are presented in column (2). The monetary premia are listed in column (5). Now suppose the subject faces the same decision, but rounds the filled-in LL amounts to the next whole dollar. These LL amounts are in listed in column (6). Column (7) displays the discount rates that are implied by these choices.

Table 2: The Effects of Rounding

Amount today (1)	Future Amount (2)	Horizon (3)	AER (4)	Monetary premium (5)	Future Amount (6)	AER (7)
$5	$5.04	28	10%	$0.04	$6	260.71%
$5	$5.08	56	10%	$0.08	$6	130.36%
$5	$5.12	84	10%	$0.12	$6	86.90%
$5	$5.15	112	10%	$0.15	$6	65.18%
$10	$10.08	28	10%	$0.08	$11	130.36%
$10	$10.15	56	10%	$0.15	$11	65.18%
$10	$10.23	84	10%	$0.23	$11	43.45%
$10	$10.31	112	10%	$0.31	$11	32.59%
$15	$15.12	28	10%	$0.12	$16	86.90%
$15	$15.23	56	10%	$0.23	$16	43.45%
$15	$15.35	84	10%	$0.35	$16	28.97%
$15	$15.46	112	10%	$0.46	$16	21.73%
$20	$20.15	28	10%	$0.15	$21	65.18%
$20	$20.31	56	10%	$0.31	$21	32.59%
$20	$20.46	84	10%	$0.46	$21	21.73%
$20	$20.61	112	10%	$0.61	$21	16.29%
$25	$25.19	28	10%	$0.19	$26	52.14%
$25	$25.38	56	10%	$0.38	$26	26.07%
$25	$25.58	84	10%	$0.58	$26	17.38%
$25	$25.77	112	10%	$0.77	$26	13.04%
$30	$30.23	28	10%	$0.23	$31	43.45%
$30	$30.46	56	10%	$0.46	$31	21.73%
$30	$30.69	84	10%	$0.69	$31	14.48%
$30	$30.92	112	10%	$0.92	$31	10.86%

As one compares the implied discount rates after rounding for choices with a given principal and varying horizons, it becomes apparent that they decline with the horizon. For example, with a principal of $5, choices with rounding imply discount rates that fall from 260.71% to 65.18%, as the horizon increases from 28 days to 112 days. Without their proper attribution to rounding behavior, the declining discount rates would appear to be due to hyperbolic time preferences.

Furthermore, the implied discount rates, after rounding and for a given horizon, exhibit what appears to be a "magnitude effect." For example, a comparison of the LL amounts across options with a 28 day horizon reveals that the implied discount rate decreases from 260.71% to 43.45% as the principal increases from $5 to $30.

Rounding behavior may occur when people compare *similar* quantities. The concept of similarity is of central importance in cognitive psychology. Various cognitive processes, such as categorization, problem solving, norm foundation and skill transfer, are based on similarity perception. A particularly influential approach to similarity is Tversky's (1977) feature-based approach. This approach represents objects by assuming sets of features associated with them. The similarity of two objects is defined as a metric that depends on three arguments: the intersection of the respective feature sets of both objects, and the two relative complement sets. Tversky's model expresses similarity relations between objects as a linear combination of the measures of their common distinctive features. The weights associated with the components of this combination reflect the salience or prominence of the various features (Tversky 1977, 329-333).

Tversky's work influenced Rubinstein's approach (1988), which emphasizes the role of similarity in decision making under risk. Rubinstein (1988) suggests that individuals comparing two lotteries first look for dominance relations. In case no lottery dominates the other, they go on to assess the similarity or dissimilarity of the features of the lotteries, i.e. prizes and their associated probabilities. If the lotteries are similar with respect to one feature but dissimilar to the other one, then the dissimilar feature is the decisive factor. For example, consider the lotteries $L_1 = (\$4000, p = 0.2; \$0, p = 0.8)$ and $L_2 = (\$3000, p = 0.25; \$0, p = 0.75)$. The expected value (EV) of L_1 is greater. However, even if we observe a choice of L_1, it would not necessarily have to be due to its higher EV. Standard economic theory could explain this choice with a high level of risk aversion or probability pessimism or a combination of both. A similarity based approach offers a different account. The probabilities of winning a positive amount may arguably be perceived as similar, whereas most people would regard the positive dollar amounts as dissimilar. In this case the dollar amounts become the decisive feature and most people will prefer L_1.

Leland (1994, 151-153) follows Rubinstein's (1988) approach, but proposes some modifications. First, he suggests that people initially attempt to decide according their true preferences. Leland (1994) assumes that preferences are assessed in accordance with the axioms of EUT. Only if differences in expected utility are sufficiently large can individuals determine a clear preference relation and the decision making process is complete. Otherwise, Leland (1994) proposes a second stage in which prizes and probabilities are compared in terms of *equality*. If this step remains inconclusive, a comparison of prizes and probabilities in terms of *similarity* is performed.

Leland's (1994) approach emphasizes that the purpose of resorting to equality or similarity assessments is to reach a decision in situations in which optimization behavior according to EUT would be associated with costs in terms of cognitive effort. The notion that the subjective utilities of similar rewards are difficult to discriminate may be illustrated by an analogy to the early work in psychophysics. Early researchers in this area, notably Weber (1834) and Fechner (1860), were interested in the perception of physical stimuli and the ability of people to discriminate between similar stimuli. For example, suppose the stimulus of interest is weight and that two light objects, weighing 15 oz and 20 oz, are being compared. Due to the relatively small weight difference of 5 oz, people will find it difficult to judge which object is heavier and are likely to make mistakes. However, if the weight of the heavier object is increased to 25 oz, the probability of making errors will decrease. If the weights of both objects are increased by one pound, judging which object is heavier, the object weighing 1 lb and 15 oz or the object weighing 1 lb and 25 oz, will likely become even more difficult and the error rate is likely to increase (Loomes 2005, 302-303). It seems plausible that comparing similar risky payoffs is even more cognitively demanding than comparing similar physical stimuli.

The difficulties of forming unambiguous preference relations have long been acknowledged by stochastic choice models. An early example from the economic literature is Georgescu-Roegen's (1958) "threshold model." The analogy to psychophysics illuminates how the cognitive effort of reward evaluation may prevent optimizing choice behavior: the difficulty lies in judging how desirable receiving a monetary amount is compared to receiving a similar, only slightly larger amount, in the future. In order to reduce the cognitive burden of finding the optimal choice in time preference tasks, subjects may perform the mental operation of rounding LL amounts to the nearest whole dollar. This has concrete design implications.

The LL amounts are similar when AER of a typical magnitude are applied to small amounts. Column (5) of Table 2 displays monetary premia that remain tiny as long the principals are small. This illustrates that rounding is more likely to occur when the principals are small. Therefore, the stakes of the present experiment are not limited to small amounts. Small principals of $10 and $30 are supplemented by a moderate principle of $60 and rather sizeable principals of $100 and $300. Budgetary constraints make it necessary to use only a 10% chance that the two large amounts and their corresponding LL amounts are paid. The other amounts are paid with certainty, if selected at the end of the time preference task.

4. Reward bundling

Impulse control has been a topic of great interest among scholars and clinical practitioners. Ainslie (1992, 130-144; 2001, 74-85) summarizes this literature and identifies four general strategies frequently recommended to make people self-commit to advantageous future choices. Those suggested strategies are the use of extrapsychic mechanisms that make impulsive choice physically impossible, the manipulation of attention so that the source of temptation is ignored, and the preparation of emotion that cultivates adverse responses to temptations. The fourth tactic, which is probably most often recommended, is the exercise of *willpower*. Referring to this human faculty is intuitively appealing. After all, most people have some notion of making resolutions and trying to resist temptations. However, the exact nature of the *will* is often left vague and attempts to give precise definitions are rare.

Ainslie (1992, 143-144; 2001, 78-84) relates the notion of the will to the application of personal rules. People often experience the belief that they exercise willpower when their current decision is approached as if it was guided by a universal principle. Ainslie (2001, 80-81) points out that many authors from different backgrounds have related impulse-control to choosing in broad categories or choice guided by universal rules. If people actually choose following this guideline, present decisions are approached as one instance in a series of related decisions. Consistent patterns of choice emerge, which are even observed in experiments with animal subjects. Human subjects often experience such consistent choice behavior as a "matter of principle," the result of "character" or "resolve," or the application of a personal rule.

The effect of personal rules on behavior is not derived from consistency *per se*, but from the mechanism of *reward bundling*. Suppose that an individual prefers the SS option in an isolated SS/LL decisions. However, when this decision is only the first of a whole series of decisions, the choices may be added together and then evaluated, rather than being evaluated in isolation. In other words, the individual *bundles* the rewards. The aggregated LL rewards may yield a higher utility than the aggregated SS rewards. In this case the subject will choose LL, even if SS rewards are preferred when the decision is made in isolation.

Ainslie (1992, 145; 2001, 81) assumes *additivity*, such that the aggregate utility of a series of rewards equals the sum of their individual utilities. Moreover, his argument hinges on the hyperbolic form of the underlying discount functions. Hyperbolic discounting implies a steep slope as rewards are delayed, and a shape that flattens as the time horizon increases. Hyperbolic discount curves are steeper than exponential discount functions (D^E) when delays are short, but flatter than both exponential and quasi-hyperbolic (D^{QH}) discount functions, when delays are long. The high tails of hyperbolic curves suggest greater effects from aggregation than one would expect from theD^{QH}model (Ainslie 2012, 12-17).

A simple example, based on the discussion in Ainslie (2012, 13) is illustrated in Table 3. Table 3 displays a hypothetical situation in which an individual faces ten decisions that involve a series of SS rewards of $100 and LL rewards of $103. All ten LL rewards are available thirty days after the corresponding SS rewards. The first SS is available immediately and every subsequent SS reward is available forty days after the previous one. Three models are considered, for illustrative purposes: the exponential, the quasi-hyperbolic, and the hyperbolic model. As a simplification linear utility functions are assumed. As the discussion in subsection 2.2.1 showed, this assumption is not realistic and can lead to seriously biased estimates of discount rates when it is imposed on actual choice data. However, assuming risk neutrality for simplification does not affect the argument in the present example, which merely serves illustrative purposes. The exponential model in panel A uses the specification from (1a) and assumes $\delta = 0.5$. Also considered is the quasi-hyperbolic model, with the specifications from (6a) and (6b) and $\beta = 0.7$ and $\delta = 0.5$. Finally, the hyperbolic model is considered. The specification D^{H2} from (4) is used, with $k = 0.4$.[7]

Table 3: Three Discounting Models and the Effects of Reward Bundling

SS Delay (days)	SS (amount)	SS (discounted)	LL Delay (days)	LL (amount)	LL (discounted)
A. Exponential Model					
SS Delay (days)	SS (amount)	SS (discounted)	LL Delay (days)	LL (amount)	LL (discounted)
0	$100.00	$100.00	30	$103.00	$99.62
40	$100.00	$95.65	70	$103.00	$95.29
80	$100.00	$91.50	110	$103.00	$91.15
120	$100.00	$87.52	150	$103.00	$87.19
160	$100.00	$83.72	190	$103.00	$83.40
200	$100.00	$80.08	230	$103.00	$79.78
240	$100.00	$76.60	270	$103.00	$76.31
280	$100.00	$73.27	310	$103.00	$72.99
320	$100.00	$70.08	350	$103.00	$69.82
360	$100.00	$67.04	390	$103.00	$66.79
Sum:		$825.45			$822.35
B. Quasi-Hyperbolic Discounting					
SS Delay (days)	SS (amount)	SS (discounted)	LL Delay (days)	LL (amount)	LL (discounted)
0	$100.00	$100.00	30	$103.00	$69.74
40	$100.00	$66.96	70	$103.00	$66.71
80	$100.00	$64.05	110	$103.00	$63.81
120	$100.00	$61.26	150	$103.00	$61.03
160	$100.00	$58.60	190	$103.00	$58.38
200	$100.00	$56.05	230	$103.00	$55.84
240	$100.00	$53.62	270	$103.00	$53.42
280	$100.00	$51.29	310	$103.00	$51.09
320	$100.00	$49.06	350	$103.00	$48.87
360	$100.00	$46.93	390	$103.00	$46.75
Sum:		$607.82			$575.64
C. Hyperbolic Discounting					
SS Delay (days)	SS (amount)	SS (discounted)	LL Delay (days)	LL (amount)	LL (discounted)
0	$100.00	$100.00	30	$103.00	$99.72
40	$100.00	$95.80	70	$103.00	$95.66
80	$100.00	$91.94	110	$103.00	$91.92
120	$100.00	$88.38	150	$103.00	$88.46
160	$100.00	$85.08	190	$103.00	$85.25
200	$100.00	$82.02	230	$103.00	$82.26
240	$100.00	$79.18	270	$103.00	$79.48
280	$100.00	$76.52	310	$103.00	$76.88
320	$100.00	$74.04	350	$103.00	$74.45
360	$100.00	$71.71	390	$103.00	$72.16
Sum:		$844.66			$846.24

[7] Note that those values are merely hypothetical. They were chosen to illustrate the logic of the reward bundling argument and are not based on actual estimates.

All three models suggest that the SS option would be preferred to the LL option if the first decision was presented in isolation. The exponential model in Panel A does not predict any change in this preference pattern when the sum of all ten rewards in the series is evaluated. For the exponential model the discount rate is constant. Hence, the preference order of the subsequent decisions is the same as in the first decision, and so is the preference order of the aggregated rewards. The quasi-hyperbolic model in Panel B does also not predict a reversal of preferences from SS to LL, as a consequence of reward aggregation. Ainslie (2012, 12) acknowledges that it is in principle possible that with a D^{QH} specification, for some parameter values, the aggregate discounted LL rewards of a series could be greater than *the aggregate rewards of the* SS alternatives. However, Ainslie (2012) points out that this can only happen when the individual already has low discount rates. Only the hyperbolic model in Panel C predicts that a person with relatively high discount rates may overcome his initial preference for the SS reward when he perceives his choices as being between whole series of LL versus SS rewards. The difference between the two sums of rewards decreases with the length of the series, and the aggregated LL rewards become ultimately more attractive.

An obvious threat for reward bundling as a strategy for impulse control comes from the possibility of "exceptions." Personal rules can be defined and re-defined in various ways, so that many or even most situations are exceptions from the rule. An example is the alcoholic, who vows not to drink but then declares every possible occasion an exception, such as a social event, the holiday season, a new brand of liquor, the cold weather, or the tough day he had. Indeed, any possible explanation is employed, but the acceptance of being an alcoholic is avoided. The incentive to find exception is apparent from Panel C in Table 3. It is obvious that the total payoff could be increased by choosing SS in the first decision and then committing to choose LL in subsequent decisions. Of course, the same incentive is present when the remaining decisions are approached. Obviously, such exceptions from the general commitment to choose LL threaten to make impulse control collapse completely. One can envision this as a situation in which short-term interests constantly threaten to undermine long-term interests.

Ainslie (1992, 155-173; 2001, 90-100) conceptualizes such “limited warfare” among successive motivational states as a repeated *Prisoner’s Dilemma*. An individual’s successive motivational states can either *cooperate* or *defect*, by choosing in favor of the commonly shared long-term interest or giving in to the temptation of peculiar short-term interest, respectively. Each individual decision may come with the threat of *defection*. However, once a series of such defections has occurred, cooperation may collapse and each motivational state will only choose SS. On the other hand, if several subsequent motivational states choose to cooperate, the result may be a tacit agreement to cooperate, which the individual experiences as a personal rule. Such a personal rule may emerge when the current choice is perceived as predictive for future choices. This precedence function of the current choice may provide sufficient incentive for the current motivational state to cooperate. Clear guidelines that clarify what constitutes cooperation in this intrapsychic bargaining situation serve to limit the availability of exceptions, thus supporting the adherence to personal rules. A well-known example is the Alcoholics Anonymous policy to discourage attempts at “controlled drinking.” Instead, they promote strict abstinence for their members, a rule which provides a *bright line* and leaves no room for exceptions.

The literature in economics has produced several comparable models of strategic interaction among motivational states. A notable example that can be related to Ainslie’s bundling approach is Fudenberg’s and Levine’s (2006) model. Their *dual-self* approach models intertemporal choice as the result of the interaction among a patient *long-run self* and a sequence of *short-run selves*. The long-run self exercises costly impulse-control by constraining the choices of the myopic short-run selves. In Fudenberg's and Levine's (2012) updated model, the short-run selves are not completely myopic and the long-run selves impose their constraints to maximize the discounted utilities of the short-run selves. The notion of a dual self can be interpreted in psychological terms, with a dual cognitive process informing the individual's decision making process.

Psychologists have contributed a rich literature on *dual process* models. The notion of dual cognitive processes has been applied to several areas in reasoning, social cognition, as well as judgment and decision making. The various models have in common that they assume that manifest behavior is determined by two modes of processing: an automatic process that is often associated with fast, unconscious, heuristic-driven, and impulsive cognition and a controlled process that is often associated with slow, conscious, deliberate, and sophisticated processing (Barrett et.al. 2004, Evans 2008).

These different processes do not necessarily have to be reduced to a one-dimensional utility index, as most economic models suggest. An alternative strategy is exemplified by Lopes' (1987) SP/A model – a psychological *dual criterion* approach to decision making under risk. The model proposes two criteria that people consult, when evaluating lotteries. The security-potential (SP) criterion, implies a decumulative probability weighing model, much like rank-dependent utility (RDU) in economics (Quiggin 1982). The aspiration (A) criterion, on the other hand, introduces a reference point (Lopes & Oden 1995, 291). Both criteria are employed independently but simultaneously.

The principles that guide the analysis of such a dual criteria model are instructive for the experiment that is reported in this thesis. Economic models do in principle allow for multiple decision processes. However, traditional economic models usually integrate those latent processes into a one-dimensional criterion. This is why they do not necessarily help understand how observed choices can result from the interplay of impulsive and controlled processes. By contrast, dual criteria models provide such a guideline. They allow for the possibility of multiple criteria for multiple decision processes (Andersen, Harrison, Lau, & Rutström 2014a). The relative contribution of each process to the observable outcome can be estimated by a finite mixture model using maximum likelihood. This approach estimates the grand likelihood of the dual criteria model as the probability weighted average of the conditional likelihood estimates of each criterion being correct (Harrison & Rutström 2009; Andersen, Harrison, Lau, & Rutström 2014a).

Ainslie (2012, 11) argues that “[w]e are born with steep discount curves, and continue to show them when the reward is an increase in our immediate comfort …We must acquire by learning whatever devices let us achieve the banker-like preferences we sometimes show." Ainslie (2012) claims, in other words, that people are born with hyperbolic preferences but learn to overcome them. When they bundle rewards their behavior appears more patient, to the extent that their observable choices are actually better described by exponential discounting (Ainslie 2001, 100-104). This is consistent with empirical results that show that the prevalence of exponential discounting is much higher than often assumed (Andersen, Harrison, Lau, and Rutström 2014b). Ainslie (2012) acknowledges this evidence but regards it as the result of the bargaining process, rather than evidence against it.

These theoretical considerations have important implications for experimental design and analysis. If two latent processes determine intertemporal choice, one impulsive and one controlled, the analysis can be approached in a way that is equivalent to the estimation of dual criteria models. A finite mixture model can then be used to estimate the conditional likelihood of each process (Coller, Harrison, and Rutström 2012).

Earlier experimental work on reward bundling includes a study by Kirby and Guastello (2001). Their experiment 1 uses salient monetary rewards and "pseudo-volunteers." That is, their subject pool includes students, whose participation was rewarded with extra-credit for an introductory psychology class. They implement different treatments by informing subjects that they will be contacted again to choose and that their current choice imposes varying constraints on future choices. The experiment reported in this thesis uses a simpler design. The theory of reward bundling is concerned with behavior that involves decisions about the future. This does not necessarily mean that those decisions have to be made in the future. Nevertheless, Kirby's and Guastello's (2001) work is relevant for the present design, since it inspired treatment conditions that are implemented.

Four experimental treatments T4 to T7 are added, which are, henceforth, referred to as *bundling treatments*. Reward bundling predicts that subjects are more likely to choose the LL options when they are presented with a series of choices. This may be the case, even if the subjects would have preferred the SS options, had they been presented in isolation. Treatment T5 was introduced to test this prediction. Participants were presented with two pairs of SS/LL rewards. Their choice in one pair determines the choice in the other pair. It is not possible to choose the SS option in one pair and the LL option in the other pair. That is, if a subject chooses one option, the experimental design forces him to "choose" the equivalent option in the other pair. This *forced condition* is equivalent to the "imposed condition" that Kirby and Guastello (2001) used in their design.

As explained above, a subject who is not "forced" to choose in a way that is consistent with the other decisions in a series may seek to maximize his payoff by choosing SS in the first decision and trying to commit to choosing LL in the remaining decision. In this case, the first choice would be regarded as an exception, with the intent to apply self-control in the remaining decisions. The obvious problem is that the same reasoning may apply to subsequent choices as well, with the effect that impulse-control collapses. Ainslie (2012, 17-19) predicts that human subjects will show a behavioral tendency to perceive a current choice between SS and LL rewards in a series of decisions as a test case for their remaining choices. Consequently, they will be more apt to prefer the LL reward when they face a series of SS/LL decisions, rather than when they see the pair of alternatives as an isolated choice. Kirby and Guastello (2001) call this the "precedent effect." It is not possible to decide whether such subjects apply this strategy if subjects are "forced" to choose consistently in all choice pairs. This is why another treatment T4 is introduced. This treatment also includes two decisions. However, treatment T4 implements a *free condition*, in which no choice constraints are imposed.

Whether subjects bundle rewards and choose the LL option in both pairs depends only on their perceived linkage between options. An increased tendency to choose LL is only one indicator for the presence of reward bundling. In addition, the perceived linkage of decisions, which is implied by the "precedent effect," should be reflected by a seemingly increased payoff *scale*. The *scale* indicates the effective magnitude of the payoffs. It equals 1 in most *control treatments*. Only in treatment T2, where the payoffs are tripled, does the scale equal 3. The *forced condition* in treatments T5 implies a scale of 2. These scale values are due to the fact that every choice in the *forced condition* leads to twice the nominal payoff value in T5. If reward bundling is effective as a mechanism of impulse control, one should see similar tendencies in the *free condition*. That is, subjects who bundled rewards would display choice behavior that looks like the scale of the payoffs was multiplied by the number of decisions. For example, there are two *free* decisions in treatment T4. Reward bundling implies choice behavior that makes the payoff scale appear as if it has a value of 2.

Similarly, reward bundling suggests that immediately available payoffs resemble choices with a FED, since the temptation of immediacy is successfully overcome. This is why one would expect that *first* choices in bundling treatments look like choices from treatments where a FED was present.

On a structural level, the analysis will focus on the effects of treatments T4 and T5 on the mixture probability of the latent decision processes. The control treatments, T0 to T3, make it possible to infer the likelihood of the choice parameters, conditional on the choice discount specification being true. Ainslie's theory (2012) predicts the presence of exponential as well as hyperbolic discounting. However, if reward bundling is effective, one would expect to see an increase in the mixture probability of exponential discounting in treatments T4 and T5.

To restate the theory, reward bundling involves the integration of a series of future rewards by summing their separate present values. Impulse-control is successful if the sum of the series of LL rewards exceeds the sum of the series of SS rewards, even when the SS rewards are preferred in the first pair of the series. One would expect that extending the series of decisions should increase the tendency to bundle choices. To test the effect of an additional decision, two additional treatments are added. Treatment T6 is again a *free condition*. It is equivalent to treatment T4, in the sense that no constraints are imposed on the choices. However, treatment T6 adds a third decision, so that subjects face a decision triple.

Finally, treatment T7 implements another *forced condition*. It is equivalent to treatment T5, since subjects must choose consistently and cannot choose SS in some and LL in other decisions. Treatment T7 presents subjects with three decisions that are linked in this way. The descriptive and structural evidence for reward bundling in the three-decision treatments T6 and T7 should be at least as strong as in the two-decision treatments T4 and T5.

5. Descriptive Methods

This chapter provides a brief discussion of the descriptive econometric methods used for a first exploration of some basic hypotheses derived from the theory of reward bundling. For this purpose it will be useful to restate some fundamental predictions that follow from Ainslie's theory (1975, 1992, 2001, 2012):

- When presented with isolated SS/LL decisions, subjects will exhibit choice behavior that is best described by hyperbolic discounting, as specified by D^{H1} and D^{H2}.
- When subjects choose between a series of SS/LL decisions at once, they will display a greater tendency to choose the LL rewards than subjects who face the same decisions one at a time.

An in-depth analysis of underlying choice behavior will be accomplished by structural estimation. Until then, a descriptive examination is intended to shed light on general tendencies that can be found in the choice data. The two methods discussed in this chapter explain the techniques used to evaluate whether the observed choices are consistent with the above predictions.

5.1 Local polynomial regression

The bundling predictions imply that subjects must face decisions that involve isolated SS/LL pairs and decisions that involve series of SS/LL pairs. This suggests that the whole sample must be divided into several subsets, where each subset is subjected to different treatment conditions. One can broadly distinguish between *control treatments*, in which subjects have to choose between isolated SS/LL pairs, and *bundling treatment*s, in which the choice task requires that subjects choose between whole series of SS/LL pairs. One can then investigate whether the fraction of LL choices in *bundling treatments* was greater than in the *control treatment* groups. Such general choice tendencies would constitute credible first evidence consistent with the presence of reward bundling.

Until the observed choice behavior is linked to structural models, the analysis remains agnostic, regarding the functional form of the data generating process. It is appropriate that the initial analysis relies on only a minimum of parametric assumptions. Econometricians have developed a rich set of nonparametric estimation techniques for such situations.

Consider the simple case of an outcome variable Y that, for the sake of implicitly, is explained terms of a single explanatory variable X (the general logic can easily be extended to multiple explanatory variables). A sample is drawn that consist of a set of bivariate data $\{(x_1, y_1), (x_2, y_2), \ldots, (x_n, y_n)\}$. The goal is to estimate the conditional expectation function $m(x_0) = E(Y|X = x_0)$, without assumptions about the functional form $m(\cdot)$. Nonparametric estimation rests on the assumptions that data points in the neighborhood of a given x_0 carry more information about $m(x_0)$ than data points far from x_0. Thus, it seems intuitive to approximate the function $m(x_0)$ by obtaining the local average within a predefined neighborhood of x_0. The performance of this hypothetical estimator can often be improved by using a weighted average.

These weights are assigned by a function k, called a *Kernel function*. Kernel functions are usually symmetric probability densities, such as the normal or a symmetric beta distribution.[8] Moreover, let h denote a non-negative number, the bandwidth, which defines the size of the neighborhood of x_0. The function $m(x_0)$ is estimated by

$$\widehat{m_{NW}} = \frac{\sum_{i=1}^{n} k_h(X_i - x_o) y_i}{\sum_{i=1}^{n} k_h(X_i - x_o)}. \tag{9}$$

This is the Nadaraya-Watson estimator, an early contribution to non-parametric estimation that has influenced comparable estimators, such as the Gasser-Müller estimator (Nadaraya 1964, Watson 1964, Gasser & Müller 1979). However, these estimators often exhibit undesirable statistical properties, such as bias or large variances under specific circumstances (Fan & Gijbels 1996, 13-18). Local polynomial regression retains the logic of local smoothing, while avoiding many of these problems.

Local polynomial regression involves fitting the response to a polynomial form of the regressor via locally weighted least squares. The result is a smooth function over the support of the data, which approximates the conditional expectation function $m(x_0)$. Suppose the $(p+1)^{th}$ derivative of $m(x)$ at the point x_0 exists. The unknown conditional expectation function is then approximated by a polynomial of order p. A Taylor expansion gives for x in the neighborhood of x_0 gives

$$m(x) \approx m(x_0) + m'(x_0)(x - x_0) + \frac{m''(x_0)}{2!}(x - x_0)^2 + \cdots + \frac{m''(x_0)}{p!}(x - x_0)^p. \tag{10}$$

This polynomial is fitted locally by the weighted least square problem

$$\min \sum_{i=1}^{n} \{Y_i - \sum_{j=1}^{p} \beta_j (X_i - x_0)^j\}^2 \, k_h(X_i - x_0), \tag{11}$$

[8] Since these are basically parametric assumptions, we often refer to semi-parametric modeling.

where k is again a Kernel function assigning weights to each data point and h is the bandwidth parameter determining the size of the local neighborhood. Specific applications require the choice of the degree of the polynomial, the type of the Kernel function, and especially the size of the bandwidth. If the latter is too large the model will contain large bias, whereas the estimates will be noisy if the bandwidth is too small. The solution vector of this least square problem is given by

$$\hat{\beta} = (X^T W X)^{-1} X^T W y \; , \tag{12}$$

where W is the matrix of weights. The statistical properties of the local polynomial regression estimator are discussed by Fan and Gijebels (1996, 57-107). In the software package *Stata* this technique is implemented by the *lpoly* command.

5.2 Descriptive probit models

A second descriptive technique is the use of probit models, which examine the effects of the treatments on the probability of choosing LL. The focus on choice probabilities offers a quantitative measure and takes the analysis beyond the mere examination the fraction of LL choices that is accomplished by local polynomial regression. However, the exploration remains at the descriptive level and cannot replace the thorough structural analysis that addresses the underlying choice processes in detail.

A few general remarks about discrete choice models are in order, since such models are used in both the descriptive and the structural analysis. A generic form of discrete choice models is given by

$$y = h(x, \varepsilon), \tag{13}$$

where y denotes some discrete outcome, in the present experiment the binary SS/LL choice. The process leading from the explanatory variables to the outcome is described by the function $h(\cdot)$. The decision maker's choice is completely determined by the observable covariates x and unobservable factors ε. Although the underlying process is assumed to be deterministic, ε is unobserved and the outcomes cannot be predicted exactly. Instead, a distribution $f(\varepsilon)$ is assumed and the choice probability $P(y|x)$ is derived (Train 2009, 3).

The indicator function

$$I[hx|\varepsilon) = y] \tag{14}$$

takes a value of 1 when the statement specified in brackets is true and a value of 0, otherwise. In the present experiment, the indicator function takes a value of 1 if LL is chosen and a value 0 if SS is chosen. Using the language of indicator functions, the choice probability can be restated as the expected value of the indicator function, where the expectation is over all possible values of the unobserved factors, i.e.

$$\begin{aligned} P(y|x) &= P(I[h(x|\varepsilon) = y] = 1) \\ &= \int I[h(x|\varepsilon) = y] f(\varepsilon) d\varepsilon. \end{aligned} \tag{15}$$

Various types of discrete choice models can be distinguished by how the density $f(\varepsilon)$, the distribution of the unobserved portion of utility, is specified and how the integration in (15) is performed (Train 2009, 3-7). The probit model is derived from the assumption that $f(\varepsilon)$ is univariate normal and the integration is performed numerically through simulation.

Specifications of $h(\cdot)$ formalize decision models that describe the process of converting explanatory variables into observable choice behavior. Specific discounting processes are considered in the following chapter. The probit model can also be used to relate personal and choice characteristics to behavior, without explaining how exactly the decision is made. This approach is explained in this section, which is why the probit models are called *descriptive*. Descriptive probit models extend the analysis beyond polynomial regressions from the previous section, since they make it possible to condition choice probabilities on observable characteristics of the individual decision maker or the choice situation.

Once the observable covariates are included in the probit model, the issue of understanding its results arises. The interpretation of the coefficients in linear regression models is straightforward. Given a linear model $y = x_1\beta_1 + \cdots + x_k\beta_k + \varepsilon$, the marginal effect of each explanatory variable is constant and given by its coefficient, i.e. $\frac{\partial y}{\partial x_k} = \beta_k$. Determining the marginal effects of explanatory variables in probit models is more involved (Cameron & Trivedi 2005, 465-474). The effect of a particular explanatory variable x_k is not constant, but dependent on the value of the other covariates. The probability that a particular outcome y_i takes a value of 1, conditional on the covariates x is given by

$$P(y_i = 1|x) = \Phi(x\beta), \quad (16)$$

Where $\Phi(\cdot)$ is the standard normal cumulative distribution function (cdf). The marginal effect of change of k-th regressor x_k on the probability of the i-th decision y_i being 1 given by

$$\frac{\partial P(y_i=1|x_i)}{\partial x_{ik}} = \phi(x_i\beta)\beta_k. \quad (17)$$

When the marginal effect of a particular independent variable x_k is estimated, one must control for the other covariates in the model. There are several ways to accomplish this. The following explanations assume, for the sake of simplicity, that x_k is a discrete binary variable. One popular technique is setting all other covariates to their mean values and assessing the difference in choice probability for both values of x_k. This method estimates the *marginal effect at the means* (MEM). However, this may result in mean values that cannot represent any actual subjects. For example, in the dataset from the present experiment, the variable BLACK has a mean value of

0.57 and the mean value of FEMALE is 0.58. A more intuitive approach, which is employed in this section, is calculating the *average marginal effect* (AME). This approach considers each individual record, leaves all other covariate values as they are, and only varies the variable x_k, whose marginal effect is evaluated. The probit model is then used to estimate the choice probabilities for both values of x_k. The difference in choice probabilities is the marginal effect for this particular case. This process is repeated for each observation so that, when all individual marginal effects are found, the AME can be calculated. The statistical package *Stata* estimates the AME using the *margins* command.

6. Structural Estimation

This chapter presents the principles of structural data analysis. Structural models relate theories of choice, which are often deterministic, with actual observable choice behavior. A challenge is posed by the fact that decision models specify processes that involve theoretical constructs that are not directly observable. Key parameters of each theory are usually unobserved and "latent," and must be inferred from behavior. Maximum likelihood estimation (MLE) is performed to assess which parameter values best explain the observed choices, conditional on the structural model of the latent decision process under consideration. The latent parameters that are of central interest are those that determine discounting behavior.

In order to infer discount parameters, one must perform joint estimation. As discussed in section 3.1, it follows from Jensen's inequality that estimating a more concave curvature of the instantaneous utility function results in lower inferred discount rates for the same observed choice data (Andersen, Harrison, Lau, and Rutström 2008). This theoretical insight has important design and econometric implications. With respect to experimental design, it suggests the use of separate tasks for the elicitation of risk attitudes and time preferences.

Econometrically, the interrelation of risk and time preferences implies their *joint estimation* via full-information maximum likelihood (FIML). Discounting parameters must be estimated, conditional on risk parameters. A consequence is that errors from the risk estimates propagate into the discounting estimates. In other words, whenever the estimates of the risk coefficients are imprecise, one cannot expect to obtain precisely estimated discounting parameters.

For the moment, it is assumed that subjects in the risk aversion experiment chose consistent with expected utility theory (EUT). This means that the probabilities of the occurrence of the outcomes enter the evaluation of the lotteries, exactly as they are induced by the experimental design, and are not transformed by decision weights. The latter option is considered below, when choice behavior consistent with rank-dependent utility (RDU) theory is assumed. The EUT model employs the following constant relative risk aversion (CRRA) specification:

$$U(x) = x^{(1-r)}/(1-r), \tag{18}$$

where x is the lottery prize and r represents a coefficient that indicates the level of constant relative risk aversion. With this specification $r = 0$ describes risk-neutrality, $r < 0$ corresponds to risk-loving preferences, and $r > 0$ corresponds to risk-averse preferences.[9]

A detailed discussion of the structural estimation of risk attitudes is presented by Harrison and Rutström (2008). When an estimate of r is considered in the ML iteration process, one can calculate the expected utility of a typical lottery i. If lottery i has j possible outcomes, its EU is given by

[9] The CRRA specification is only one possible way to parameterize risk preferences. Some alternatives, such as the expo-power function, are more flexible (Saha 1993). However, this added flexibility comes at the cost of having additional parameters to estimate. This is considered a distraction in the present chapter, which focuses on time preferences and impulse control.

$$EU_i = \sum_j p(x_j)\, U(x_j). \tag{19}$$

Then, for each decision pair an index is calculated that indicates the difference in the expected utility of both lotteries in a decision pair. Formally,

$$\Delta EU = EU_L - EU_R, \tag{20}$$

Where EU_L is the "left" lottery and EU_R is the "right" lottery in a decision pair, as they were presented to subjects. As mentioned earlier, structural estimation seeks to link theoretical constructs, which are not directly observable, to observed choice behavior. The index in (20) is such a latent variable. The connection between the latent index and the observed choices is accomplished by a *link function*. In the present case, this link function is the cumulative distribution function (cdf) of the univariate normal distribution $\Phi(\cdot)$, resulting in a probit model. Unlike the descriptive probit models in section 4.2, the structural probit models in this chapter make statements about how exactly the assumed decision process leads to the observed outcomes. The probability of choosing the "right" lottery can be written

$$prob(choose\ R) = \Phi(\Delta EU). \tag{21}$$

Thus, the latent index in (20) is linked to the observed choices by making the statement that lottery R is chosen, when the $\Delta EU > 0.5$.

This basic approach can be extended in several ways. An important addition is accounting for behavioral errors. The structural probit model cannot predict individual decision making with certainty. Decision makers may deviate from their true underlying preferences for a variety of reasons. Behavioral error specifications can account for various error sources, ranging from random deviations, due to attention lapses, to systematic violations, related to the psychology of perception and judgment. In section 3.4, in the context of the discussion of rounding and similarity, the early work in classical psychophysics was briefly mentioned. It was explained that people have difficulties discriminating between physical stimuli when they are very similar. A particularly influential behavioral error specification, due to Fechner (1860), emanates from this tradition. Its application to the evaluation of risky prospects was later popularized by Hey and Orme (1994). The inclusion of the Fechner error specification expands the latent index in (16) to

$$\Delta EU = (EU_L - EU_R)/\mu \tag{22}$$

where the new parameter μ allows the, otherwise deterministic, EUT model to account for deviations from the underlying preference structure.

Wilcox (2008) (2011) suggests an additional expansion. The intuition behind *contextual utility* originates from psychological experiments on signal detection and stimulus discrimination. These studies discovered that errors became more likely as the range of possible stimuli increased. Contextual utility follows this observation, by assuming that evaluative errors increase with the perceived range of outcomes. Econometrically, this implies that the standard deviation of the behavioral error is proportional to the range of utilities of the outcomes in a lottery pair. The contextual error specification is given by

$$\Delta EU = (EU_L - EU_R/\nu)\ /\mu, \tag{23}$$

where the new parameter ν is defined as the maximum utility over all outcomes minus the minimum utility over all outcomes in the lottery pair, i.e. over the *context* of that pair. This specification has a normalizing effect on the latent index, which remains in the unit interval. The contextual error specification is particularly parsimonious, since the parameter ν is defined, so that no additional parameter estimation is required. The specification also allows inferences, regarding “stochastically more risk averse” relationships. The latter refers to a stochastic notion of the familiar Arrow-Pratt metric of risk aversion. A *stochastically* risk averse subject is “on average” risk averse, but the metric is flexible enough to deal with choices that deviate from the subject’s general risk aversion. With the latent index remaining within the bounds of the unit interval, one can compare the stochastic risk aversion of subjects who choose in dramatically different decision contexts (i.e. who face lotteries with very different prizes).

Once the parameters of interest are defined, structural estimation can be undertaken. The log-likelihood function is

$$LL^{EUT}(r,\mu;y,X) = \sum_i \begin{bmatrix} \left(ln\Phi(\Delta \mathrm{EU}) \times \mathrm{I}(y_i = 1)\right) + \\ \left(\ln\left(1 - \Phi(\Delta EU)\right) \times I(y_i = -1)\right) \end{bmatrix}, \tag{24}$$

where the indicator function $I(\cdot)$ signifies whether the right $(y_i = 1)$ or the left $(y_i = -1)$ lottery is chosen. The parameters r and μ indicate the CRRA coefficient and the Fechner error term, respectively. The parameters can in principle be conditioned on a vector X of demographic characteristics. It is useful to constrain the parameter μ to be greater than zero.

People may not necessarily behave as if given probabilities affect their lottery evaluation with objective values. Instead, they may distort these probabilities in their perception – a process that can be described by attaching subjective weights to probabilities. Subjective weights are already discussed by Edwards (1962) and play an important role in the original version of prospect theory (PT), due to Kahneman & Tversky (1979). However, early treatments of subjective decision weights result in implausible violations of first-order stochastic dominance (FOSD). These difficulties are avoided by rank-dependent utility theory (RDU), which derives probability weights from the entire distribution over ranked outcomes, not from individual probabilities (Quiggin (1982)). The resulting decision weights reflect the subjective distortion of objective probabilities. The RDU model, which nests the previous EUT model, is considered as an alternative DGP for the risk preference data.[10] This requires the introduction of a probability weighting function. A variety of weighting functions have been proposed in the literature. The simplest specification is a power function

$$\omega(p) = p^{\gamma}, \tag{25}$$

which introduces one additional parameter γ. The log-likelihood function is then

$$LL^{RDU}(r, \gamma, \mu; y, X) = \sum_i \left[\begin{array}{c} \left(ln\Phi(\Delta \mathrm{RDU}) \times \mathrm{I}(y_i = 1) \right) + \\ \left(\ln\left(1 - \Phi(\Delta RDU)\right) \times I(y_i = -1) \right) \end{array} \right]. \tag{26}$$

Two other probability weighting functions are frequently used in applied research. One popular specification, the inverse-S shaped function, is proposed by Tversky and Kahneman (1992):

$$\omega(p) = p^{\gamma} / [p^{\gamma} + (1 - p)^{\gamma}]^{1/\gamma}. \tag{27}$$

The probability weighting function in (27) is more flexible than the power function in (22) and has additional behavioral implications. The inverse-S specification allows for a concave and a convex portion of the weighting function with a fixed point, where $\omega(p) = p$. If $\gamma < 1$, this specification reflects the often asserted empirical regularity that people overweight small probabilities and underweight large probabilities. The reverse is implied by $\gamma > 1$.

Another specification of probability weighting is contributed by Prelec (1998):

$$\omega(p) = \exp[-\eta(-\ln p)^{\phi}]. \tag{28}$$

---

[10] There are, of course, a wealth of other models of choice under risk that could have been considered. A popular alternative that does not nest EUT is PT, but several other models have been proposed by behavioral economists (Starmer 2000). Contributions from cognitive psychologists have sought to improve our understanding of the mental operations involved in the decision making process. An overview of the literature from a psychological perspective is provided by Johnson & Busemeyer (2010).

This weighting function is derived from several axioms that reflect apparent regularities of probability weighting. The specification is more flexible than the previous two, but requires the estimation of two parameters η and ϕ. There is no probability weighting when $\eta = \phi = 1$.

For the remainder of this chapter, it will be assumed that the risk preference data were generated by a process that is consistent with the RDU model with Prelec probability weighting. Results from a structural analysis of the actual choice data will be presented in the next chapter.

The estimation of discounting behavior, however, requires joint estimation of risk and time preferences. Suppose the monetary amount SS is available at time t, whereas the amount LL is delivered at time $t + \tau$. If exponential discounting holds, the present value of option SS is given by

$$PV_{SS}^{E} = (1/(1+\delta)^{t})\frac{(\upsilon+SS)^{(1-r)}}{1-r} + (1/(1+\delta)^{t+\tau})\frac{\upsilon^{(1-r)}}{1-r}, \quad (29)$$

where δ is the discount rate and υ denotes some measure of background consumption. Expression (29) says that the present value of option SS is the discounted utility of receiving the amount SS (integrated with the background consumption υ at time t and the discounted utility of receiving nothing beyond the utility of background consumption at time $t + \tau$. Utility is described using the CRRA specification from (14). Similarly, conditional on exponential discounting being the true latent process, the present value of option LL is given by

$$PV_{LL}^{E} = \left(\frac{1}{(1+\delta)^{t}}\right)\frac{\upsilon^{(1-r)}}{(1-r)} + \left(\frac{1}{(1+\delta)^{t+\tau}}\right)\frac{(LL+\upsilon)^{(1-r)}}{(1-r)} \quad (30)$$

A latent index, formally equivalent to the index considered in the risk aversion model, is defined as follows:

$$\Delta PV^{E} = \frac{PV_{LL}^{E} - PV_{SS}^{E}}{\nu}, \quad (31)$$

where the parameter ν denotes the behavioral error specification, comparable to the Fechner parameter μ in (20) and (21).[11]

The estimation maximizes the following conditional log-likelihood function:

$$LL^{E}(r,\eta,\phi,\delta,\mu,\nu;y,\upsilon,X) = \sum_i \begin{bmatrix} \left(ln\Phi(\Delta PV^{E}) \times I(y_i = 1)\right) + \\ (\ln(1-\Phi(\Delta PV^{E})) \times I(y_i = -1)) \end{bmatrix}, \qquad (32)$$

where $y_i = 1$ and $y_i = -1$ denote the choice of LL and SS, respectively. The vector Xincludes observable demographic characteristics. The joint log-likelihood function, denoted LL^{RDU_E} , can then be written as

$$LL^{RDU\_E}(r,\eta,\phi,\delta,\mu,\nu;y,\upsilon,X) = LL^{RDU} + LL^{E} \qquad (33)$$

where LL^{RDU} is the aggregate log-likelihood for risk preference choices, assuming RDU with a Prelec probability weighting function , and LL^{E} is the aggregate log-likelihood for time preference choices, assuming exponential discounting.

Structural estimation is performed in a similar fashion when alternative discounting models are studied. Consider, for example, hyperbolic discounting as specified in (4). The present value of SS is then

$$PV_{SS}^{H2} = \left(\frac{1}{1+Kt}\right)\frac{(\upsilon+SS)^{(1-r)}}{(1-r)} + \left(\frac{1}{1+K(t+\tau)}\right)\frac{\upsilon^{(1-r)}}{1-r}. \qquad (34)$$

Similarly, the present value of LL is given by

$$PV_{LL}^{H2} = \left(\frac{1}{1+Kt}\right)\frac{\upsilon^{(1-r)}}{(1-r)} + \left(\frac{1}{1+K(t+\tau)}\right)\frac{(LL+\upsilon)^{(1-r)}}{(1-r)} \qquad (35)$$

The latent index is then

$$\Delta PV^{H2} = \frac{PV_{LL}^{H2} - PV_{SS}^{H2}}{\eta}. \qquad (36)$$

The log-likelihood function of the hyperbolic discounting model is given by

[11] In the standard time preference task employed here, there is no range of uncertain outcomes involved. Hence, there is no need for contextual utility correction when behavioral errors are specified.

$$LL^{H2}(r, \eta, \phi, K, \mu, \nu; y, \upsilon, X) = \sum_i \begin{bmatrix} \left(ln\Phi(\Delta PV^{H2}) \times I(y_i = 1)\right) + \\ \left(\ln\left(1 - \Phi(\Delta PV^{H2})\right) \times I(y_i = -1)\right) \end{bmatrix}. \quad (37)$$

The joint log-likelihood function is

$$LL^{RDU\_H2}(r, \eta, \phi, K, \mu, \nu; y, \upsilon, X) = LL^{RDU} + LL^{H2}. \quad (38)$$

One can expect that impulsive and controlled processes are both present in the population. Such process heterogeneity is intuitively plausible. Moreover, it is central to the testing of the hypothesis of reward bundling, according to which deliberate processes become more prevalent when several intertemporal decisions are presented together. This coexistence of multiple decision processes implies challenges for conventional methods of data analysis, which typically incorporate the presupposition of a single data generating process (DGP). Econometrically, one can meet this challenge by employing structural *mixture models*, which explicitly allow for more than one latent DGP. Mixture models are derived from a long pedigree in the statistical literature on non-nested hypothesis testing (Harrison & Rutström (2009, 153-155), Andersen, Harrison, Lau, & Rutström (2010, 563-565), Coller, Harrison & Rutström (2012, 386-387)).

Fortunately, accounting for multiple processes requires only a straightforward extension of the structural models considered so far. The basic logic of the approach is laid out by Harrison & Rutström (2009), who apply it to decision making under risk, with EUT and prospect theory (PT) as latent DGPs. The application to structural modeling of intertemporal choice follows the same logic. Suppose there are only two DGPs. Each individual process can be estimated as outlined above, where the likelihood of each outcome is the probability of its occurrence *conditional* on the assumed DGP being true. To assess the contribution of each process, a comprehensive overarching model is then constructed. With two individual models included, this involves the estimation of only one additional *mixture parameter*.

Allowing for the coexistence of multiple DGPs raises questions regarding the interpretation of the results. A categorical approach to mixture models might associate each process with types of decision makers. This interpretation suggests that an individual employs only one specific decision process in all task domains. For instance, Harrison and Rutström (2009, 143-144) point out that they could use the mixture model to categorize each subject as being better described by EUT or PT. Similarly, one could associate each process with a specific task type (e.g. lottery choices over 2 or 3 prizes as better characterized by EUT, and lottery choices over 4 or more prizes as better characterized by PT). They reject both categorical interpretations and adopt a point of view that is shared in this thesis: people may apply multiple decision processes that are not necessarily reduced to a single decision criterion in any task (see section 2.3). The mixture probability of a specific process, for instance EUT, is then interpreted as the chance that the choice of a subject is consistent with this process. This psychological interpretation has direct implications for the present experimental results: if reward bundling is effective, one can expect the mixture probability of controlled choice behavior, as described by exponential discounting, to be higher in *bundling treatments*.

Suppose the choice data are generated by exponential *and* hyperbolic discounting behavior. Again, the specifications D^E and $D^{H2}$are used. The discounting parameter δ and K are jointly estimated with the risk parameter r from the CRRA specification in (18). Using contextual utility, the behavioral error parameter μ is added. Similarly behavioral errors in each discounting model are captured by the parameter ν.[12] All those steps are identical to the ones that

[12] It is principally possible, and not uncommon, to use only one behavioral error term for all assumed discounting process in a mixture model. However, all mixture models reported in this chapter allowed for a separate behavioral error term for each discounting process. One reason for this approach is the expectation that separate error terms might facilitate the ML estimation by enhancing numerical stability. Moreover, there were also theoretical reasons for this approach: mixture models propose two or more data generating processes and it is perfectly possible that each of these processes is influenced by different behavioral errors. For the sake of simplicity, only one error term is

one would perform if each discounting model was estimated separately. However, a mixture model requires the additional step of writing a grand log-likelihood function, which is simply the sum of the probability weighted conditional likelihoods:

$$LL^{RDU\_E-H2}(r,\eta,\phi,\delta,K,\mu,\nu;y,\omega,X) = \sum_i \ln[(\pi^{RDU\_E} \times L_i^{RDU\_E}) + ((1-\pi^{RDU\_E}) \times L_i^{RDU\_H2})] \quad (39)$$

The grand log-likelihood function LL^{RDU_E-H2} assumes that the latent discounting process is characterized in part by exponential and in part by hyperbolic discounting, and that risk preferences are consistent with RDU with a Prelec probability weighting function. It includes the new parameter π^E which is the probability of exponential discounting (with RDU consistent risk preferences) being the true model. Since only two latent processes are considered, the probability of hyperbolic discounting (with RDU consistent risk preferences) is necessarily (1-π^E).

Mixture models emerged from the statistical literature on non-nested hypothesis testing. Exponential and hyperbolic discounting are not in a nesting, in the sense that specific parameter values can transform one model into the other.[13] The quasi-hyperbolic and the Weibull models collapse to exponential discounting, if specific parameter estimates do not significantly differ from specific values. By contrast, a comparison of exponential and hyperbolic discounting cannot be accomplished by testing specific parameters. The selection between non-nested models is more complicated. The naïve comparison of log-likelihoods is generally inadequate, since log-likelihoods are influenced by the number of parameters. Some statistical selection criteria, such

mentioned during the introduction and theoretical discussion of mixture models, until the actual estimation results are reported.

[13] The extent to which two competing models are "non-nested" can be derived from the Kullback-Leibler information criterion (KLIC) (Kullback 1959). This allows classifying competing models as nested, partially non-nested, or strictly non-nested. Both non-nested hypothesis tests considered in this chapter are based on the KLIC.

as Akaike's (1973) information criterion (AIC) or Schwarz's (1978) Bayesian information criterion (BIC), do often not provide probabilistic statements regarding the model selection. For example, it is impossible to decide whether the difference in two AIC values is large or small. Nor do these measures include information about the competing model. The latter is a desirable quality for non-nested hypothesis testing.

The two non-nested hypothesis tests used in this chapter are the Vuong and the Clarke test. A common testing approach for the comparison of two nested models are likelihood-ratio tests. The strategy is to estimate the log-likelihoods of each observation in both models and to compare the ratio of both aggregate log-likelihoods (i.e. the ratio of the added log-likelihoods of all observations in one model to the added log-likelihoods of all observations in the competing model). Non-nested hypothesis tests do not use the aggregate log-likelihoods, but perform comparisons of the individual log-likelihoods of both observations. Vuong's (1989) test is derived from the Kullback-Leibler information criterion (KLIC) and proposes that one model should be selected over another if its average log-likelihood is greater. The following null-hypothesis states that both models explain the data equally well:

$$H_0: E\left[ln\frac{f(y_i|x_i,\theta^*)}{g(y_i|x_i,\gamma^*)}\right] = 0 \tag{40}$$

where $f(y_i|x_i,\theta^*)$ and $g(y_i|x_i,\gamma^*)$ are the likelihoods of the competing non-nested models f and g and θ^* and γ^* are their respective pseudo-true parameters. If the expectation in expression (36) is significantly different from zero, then one of the two models is declared preferable by this test. If $E\left[ln\frac{f(y_i|x_i,\theta^*)}{g(y_i|x_i,\gamma^*)}\right] > 0$ then model f is better. If $E\left[ln\frac{f(y_i|x_i,\theta^*)}{g(y_i|x_i,\gamma^*)}\right] < 0$, then model g explains the data better. The expectation in (40) is not actually known. However, Vuong proves that under general conditions,

$$\frac{1}{n} LR_n(\widehat{\theta_n}, \widehat{\gamma_n}) \xrightarrow{a.s} E\left[ln \frac{f(y_i|x_i, \theta^*)}{g(y_i|x_i, \gamma^*)}\right], \tag{41}$$

where $\widehat{\theta_n}$ and $\widehat{\gamma_n}$ are the maximum likelihood estimates of the parameters. Expression (41) states that $1/n$ times the likelihood ratio statistic almost surely converges to (36) as the sample size gets large. The test under H_0 then becomes

$$H_0: \frac{LR_n(\widehat{\theta_n}, \widehat{\gamma_n})}{(\sqrt{n})\widehat{\sigma_n}} \xrightarrow{D} N(0,1), \tag{42}$$

where the numerator is the difference in summed log-likelihoods and $\widehat{\sigma_n}$ is the standard error of the log-likelihood ratio. The test statistic states that if H_0 is true, and both models explain the data equally well, the average value of the log-likelihood ratio should be zero. If model f explains the data better, then the average value of the log-likelihood ratio should be significantly greater than zero. The reverse is the case if the competing model g is better. In this case, the average log-likelihood ratio should be significantly less than zero.

The Vuong test is based on the fact that, given a set of conditions, the log-likelihood ratios are asymptotically normally distributed. Clarke (2003, 2007) points out that the conditions for this asymptotic result are frequently not given. Not only are researchers often confronted with small sample sizes, the canonical correlation between competing models may also be high. This may result in distributions of log-likelihood ratios, whose distributions exhibit sharp peaks compared to a normal distribution. Clarke (2003, 2007) demonstrates that under these circumstances, the Vuong test may often fail to select the true model. He proposes a non-parametric alternative.

Clarke's (2003, 2007) test is essentially a paired sign test, applied to the differences in log-likelihoods. This nonparametric test determines whether the median log-likelihood ratio is statistically different from zero. If both models explain the data equally well, half of the individual log-likelihood ratios should be greater than zero and half should be less than zero. This suggests the following null-hypothesis:

$$H_0: Pr\left[ln\frac{f(y_i|x_i,\theta^*)}{g(y_i|x_i,\gamma^*)} > 0\right] = 0.5. \qquad (43)$$

This states that under the null hypothesis the log-likelihood ratios should be distributed around zero. If model f is better, then more than half of the log-likelihood ratios should be greater than zero. If the competing model g explains the data better, then half of the log-likelihood ratios should be less than zero. If $d_i = \ln f\left(y_i \middle| x_i\widehat{\theta_n}\right) - \ln g(y_i|x_i\widehat{\gamma_n})$, the actual test statistic can be denoted

$$B = \sum_{i=1}^{n} I(d_i|0 \leq d_i \leq +\infty)d_i, \qquad (44)$$

where the indicator function $I(\cdot)$signifies the condition that the log-likelihood differences are positive. The test statistic is distributed Binomial, with the parameters n and p=0.5.

The Vuong and the Clarke test may lead to contradictory results. Moreover, both may result in the rejection of models that are performing only slightly worse than competing models, but otherwise explain much of the observed choice behavior. However, it is not the case that only one model has to win the comparison and be considered the sole DGP. For example, various alternatives to the exponential discounting model were introduced to explain intertemporal choice behavior that appeared to be impulsive. Controlled and impulsive preferences may co-exist in a given sample, and indeed for the same subject over different choices. Hence, mixture models offer a superior testing strategy, compared to the two non-nested hypothesis tests discussed above. They reveal the extent to which the data are supported by each hypothesized DGP. Moreover, the extreme case, of one model clearly winning the comparison and explaining all data better than the competing model, can also be detected by hypothesis tests on the mixture parameter. In general, a mixture model approach will usually lead to more informative results than frequently used non-nested hypothesis tests.

7. Experimental Design and Procedure

7.1 Basic Design

The experiment consisted of eight separate sessions. In each session, a different experimental treatment was conducted. In each treatment the subjects faced decision problems where they had to choose between smaller amounts of money which would be paid sooner (SS), or larger amounts which would be paid later (LL). The task presentation followed a multiple-price list (MPL) format for the elicitation of time preferences, initially introduced for time preference elicitation by Coller and Williams (1999). User-written software, created in Visual Basic.NET by Professor J. Todd Swarthout, was used to present the task to the participants and to record their choices. In each treatment, the subjects were presented with ten decision screens, generally with different SS amounts on each screen. Moreover, every decision screen contained seven decision rows, with different implied interest rates per row. Those interest rates were not displayed to subjects. The design is summarized in Table 1.

Treatment T0 was the baseline treatment. All subsequent treatments modified specific aspects of this basic task. The SS amounts presented to subjects in different decision sheets were principals of \$10, \$30, \$60, \$100, and \$300. The amounts were presented in random order. A specific smaller amount was to be paid on the day of the experimental session, on August, 27th 2013. The LL amounts were to be paid 28 days after the session, i.e. on September, 24th 2013. The LL amounts were calculated by applying an annual interest rate from the day of the session to the day of the delayed payout. The annual interest rates ranged from 5 percent to 500 percent. It will be assumed that these are significant amounts which incentivized most subjects to carefully consider and truthfully report their preferences. A typical decision screen from T0 is presented in Figure 1.

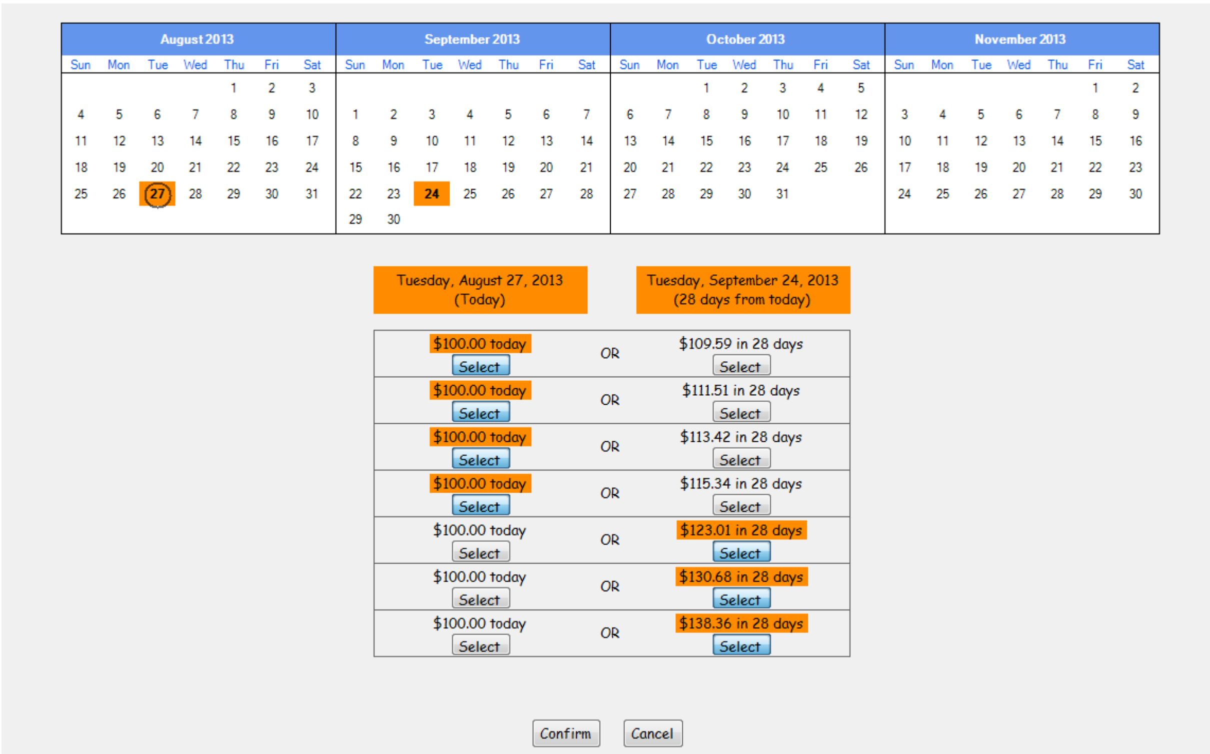

Figure 1: Decision task from treatment T0

The SS amount, $100 in Figure 1, remained constant across all decision rows on the same screen. The LL amount increased with each decision row within a given decision screen, since a higher interest rate was applied to each subsequent row. The payout dates were highlighted in the calendar on the screen, in order to provide a visual aid for the subjects' decisions. A subject faced seven decision rows on a screen and had to decide whether she preferred the SS or the LL option in each binary choice. She chose by clicking a button under the option she preferred, marked "Select." A subject who understands the task would arguably switch from SS to LL options when the LL amount was considered large enough to compensate for the delayed payout, and not switch back in subsequent rows. The hypothetical subject in Figure 1 chose the SS amount in the first four decisions and preferred the LL options from row 5 on.

Treatment T1 follows the basic design of T0, but introduced a front-end delay of one day. That is, although the session took place on August, 28th 2013, the SS amount was not paid before August, 29th. The horizon of 28 days remained constant, which resulted in September, 26th 2013 being the date when the selected LL amounts had to be paid. Figure 2 shows a typical decision screen in treatment T1.

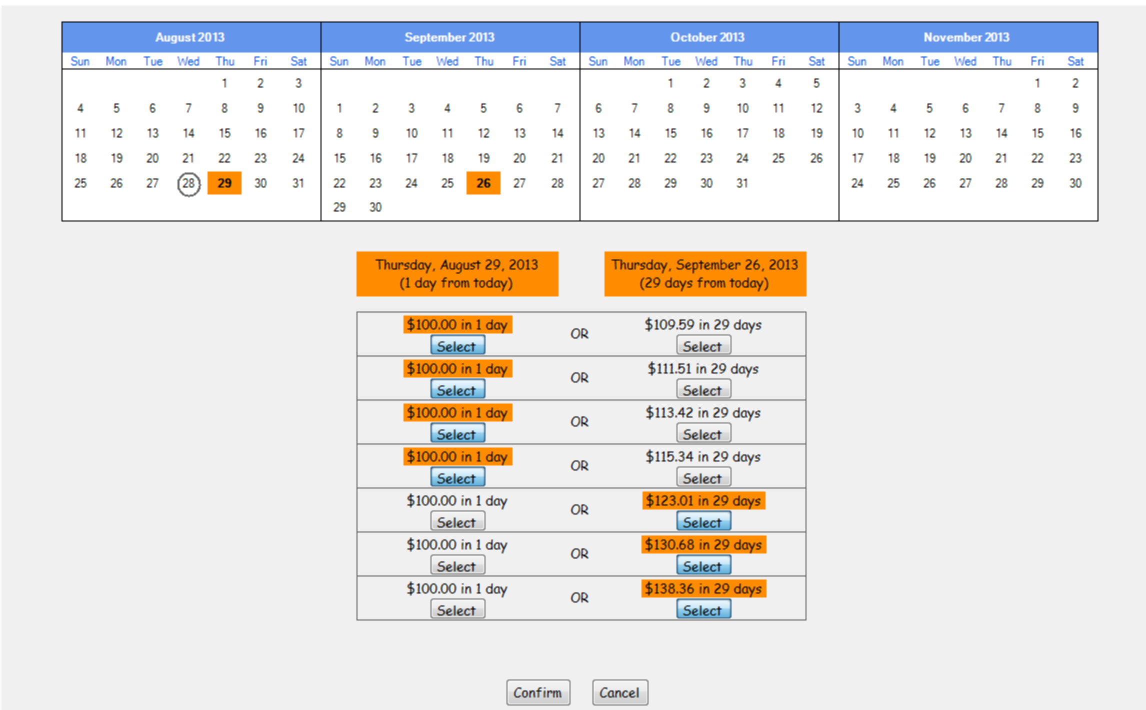

Figure 2: Decision task from treatment T1

Treatment T2 also departed from the design of T0, with the SS amounts being tripled before the usual interest rates were applied to calculate the LL amounts. Consequently, the T2 specific SS amounts ranged from $30 to $900 – a considerable increase in magnitude. Similar to treatment T0, the payout date for SS amounts was the day when the session took place – August, 29th 2013 – while the selected LL amounts were paid after 28 days – on September, 26th 2013. Figure 3 displays a typical decision screen for treatment T2.

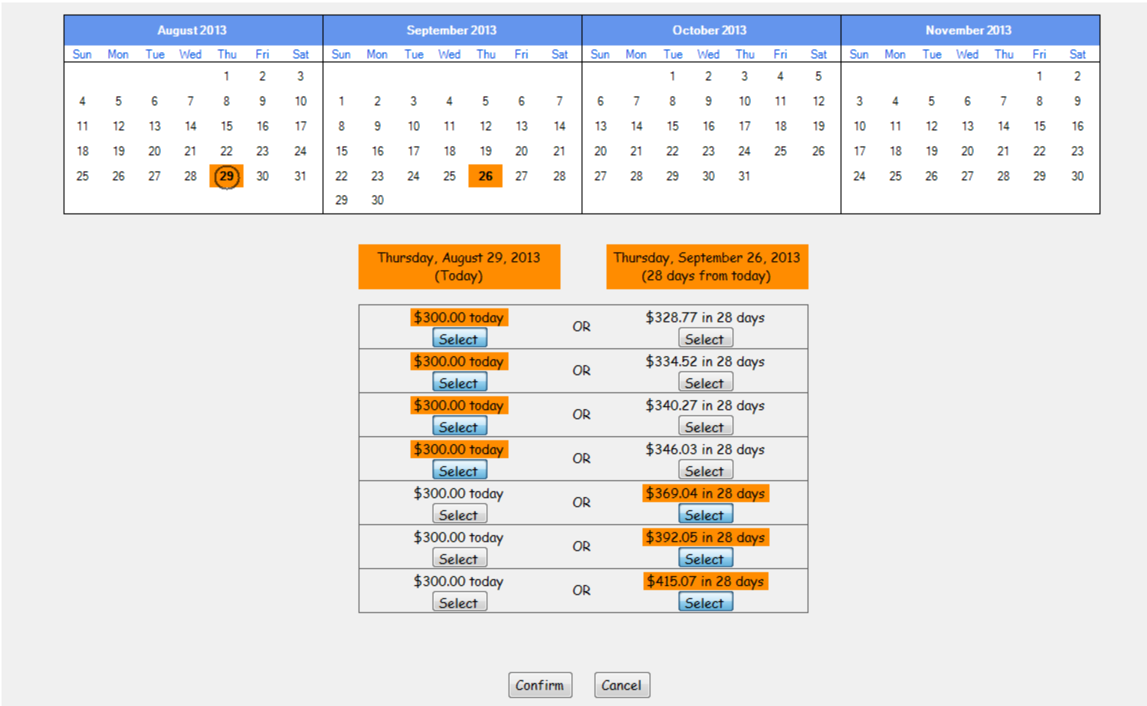

Figure 3: Decision task from treatment T2

Treatment T3 modified the baseline treatment T0 by introducing an extended FED. While the FED in treatment T1 was limited to one day, in treatment T3 it was 35 days. That is, while the session took place on August, 28th 2013 the payout date for the SS amount was October, 2nd 2013. In accordance with the constant horizon of 28 days, the payout date for LL amounts was October, 30th 2013. Figure 4 exhibits a typical decision screen in treatment T3.

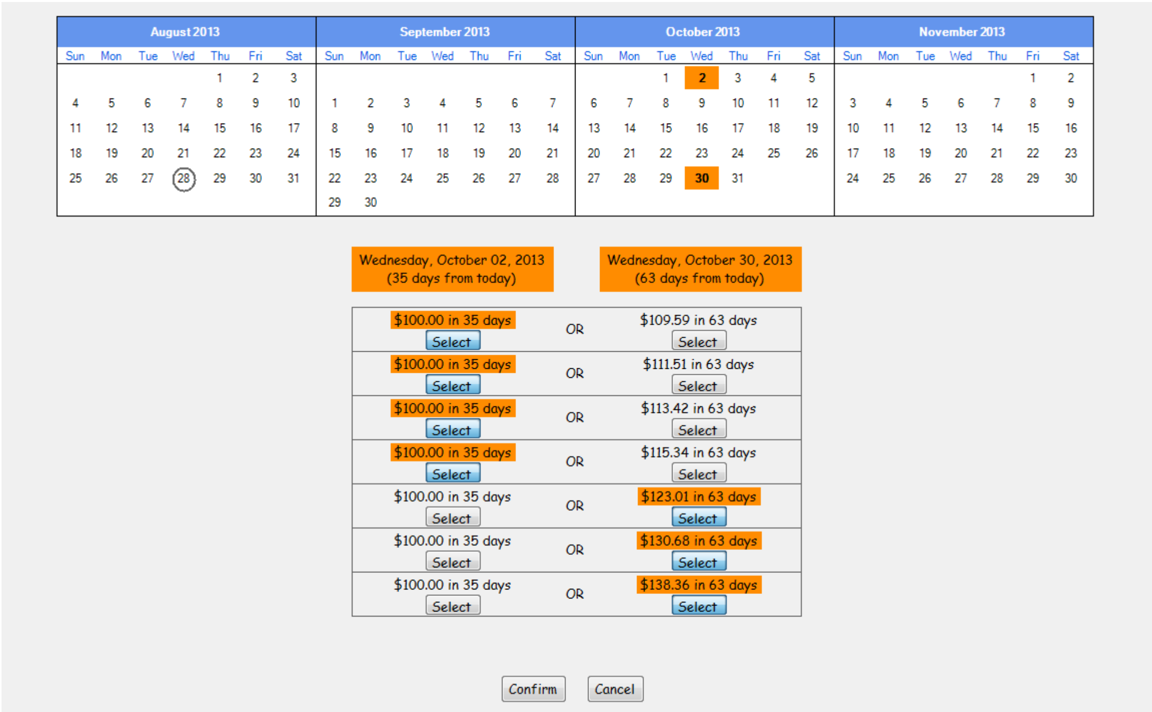

Figure 4: Decision task from treatment T3

The experimental conditions discussed so far varied the payout date or the monetary amounts involved. They are all control treatments that allow testing the effects of front-end delays and reward magnitudes on the inferred discount rates. However, since they involved only one decision per row, they do not offer evidence for bundling behavior. The central hypothesis of this thesis is that subjects, when presented simultaneously with a series of choices over time dated rewards, will have a higher tendency to choose LL options and to behave according to the exponential discounting model. The sessions that followed continued the use of ten decision screens with seven decision rows each.

Treatment T4 presented subjects with a pair of two decisions per row. Each pair involved an SS option and an LL option. The two SS amounts are identical, as are the two LL amounts in a row. However, the payout dates differ. As in treatment T0, a selected SS amount from the first decision was to be paid on the day of the experimental session, in this case on August, 30th 2013.

The payment of an LL amount from the first decision was due 28 days later, i.e. on September, 27th 2013. Just as in treatment T3, the selected SS amount from the second decision was to be paid 35 days after the session, on October, 4th 2013. As in all treatments, the horizon was constantly set to 28 days, which made November, 1st 2013 the due date for the payment of the *second* LL amount. A typical decision screen for subjects in treatment T4 is presented in Figure 5.

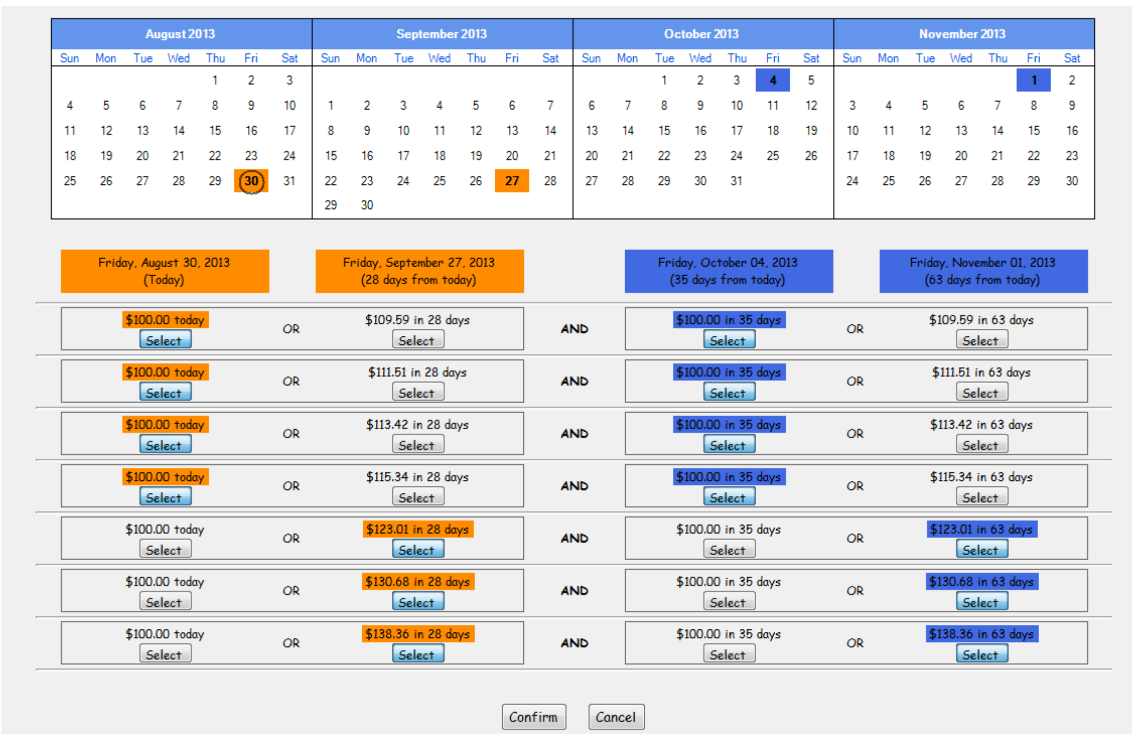

Figure 5: Decision task from treatment T4

An essential attribute of treatment T4 was that it implemented a *free choice condition*. This means that subjects were free to choose the different SS/LL options in one row across the two decisions. For example, a subject could choose the SS in the first decision and the SS amount in the second decision. Indeed, the hypothetical subject, whose choices are displayed in Figure 5, made these choices in the first row. However, she could have also chosen SS (LL) in the first decision and LL (SS) in the second decision of the first row. In other words, in the *free* *condition* of treatment T4, nothing restricted choice behavior across the two decisions so that only the same options could be chosen *within* a row. The freedom to choose different options was emphasized in the experimental instructions for treatment T4. The colors in Figure 5 do not suggest any specific choice behavior either. Different colors merely indicated different decisions in a row – they were not used to highlight the identity of amounts.

The freedom to choose different options is exactly what was constrained in treatment T5, which implemented a *forced choice condition*. Again, there were two decisions in a row. As in treatment T4, the SS amount of the first decision was to be paid on the day of the session, in this case September, 3rd 2013, whereas the corresponding LL amount was due on October, 1st 2013. In the second decision the due dates for SS and LL were October, 8th and November, 5th 2013, respectively. Once more, the two decisions in a row implemented the payment delays from treatments T0 and T3. Figure 6 shows a typical treatment T5 decision screen.

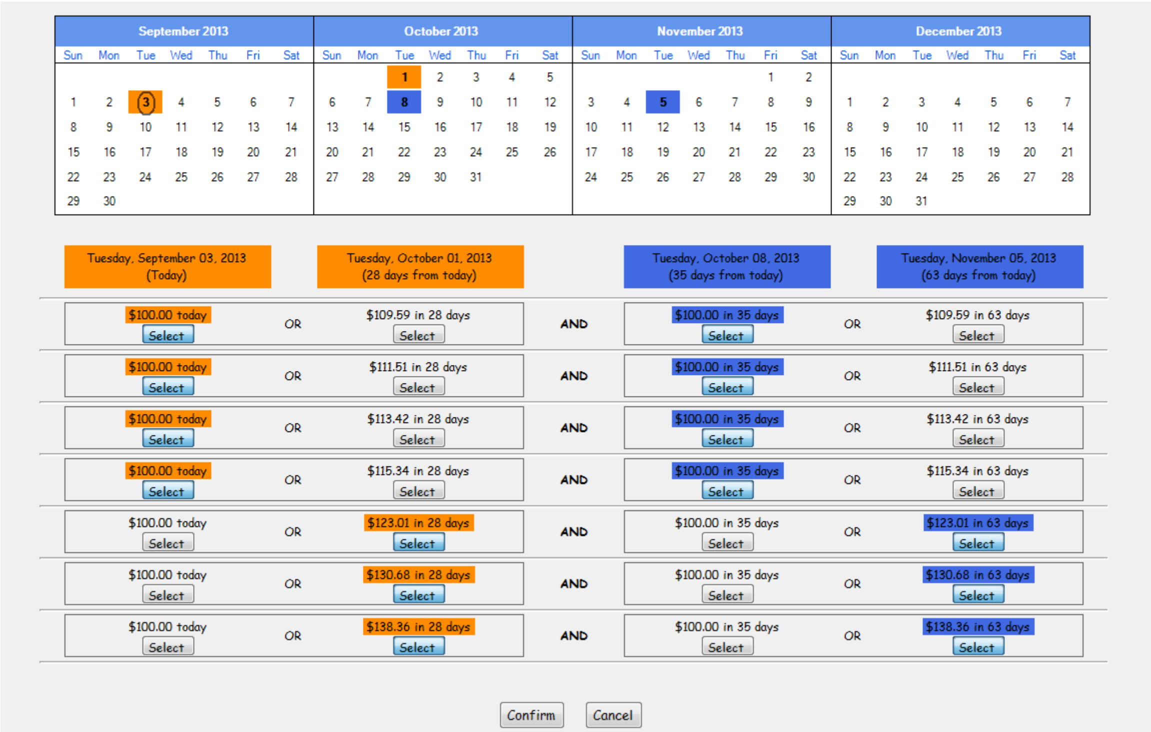

Figure 6: Decision task from treatment T5

In the *forced choice condition* the decisions in each pair were "tied together." It was not possible for a subject to choose SS in the first decision and LL in the second decision within a given row, or *vice versa*. Choosing in any one decision, by clicking the button under the option, automatically activated the choice of the corresponding decision in the same row. However, choosing options in one row did not constrain the freedom to choose different options in other rows, as long as the choices *within* a row are consistent. The experimental instructions laid out these constraints in detail.

Treatment T6 presents another *free condition*. However, this time a third decision is added to each decision row. The first decision in a triple offered the choice between an SS option, where payment was due on the date of the session and an LL option that was due 28 days later. The dates in this case were September, 4th and October, 2nd 2013, respectively. The due dates for SS and LL in the *second* decision were October, 9th and November, 6th, respectively. What distinguished treatment T6 from treatment T4 is the addition of a third decision. The SS amount for the third decision was due 70 days after the session, on November, 13th, and its LL amount was due on December, 11th 2013. A typical decision screen in treatment T6 is shown in Figure 7.

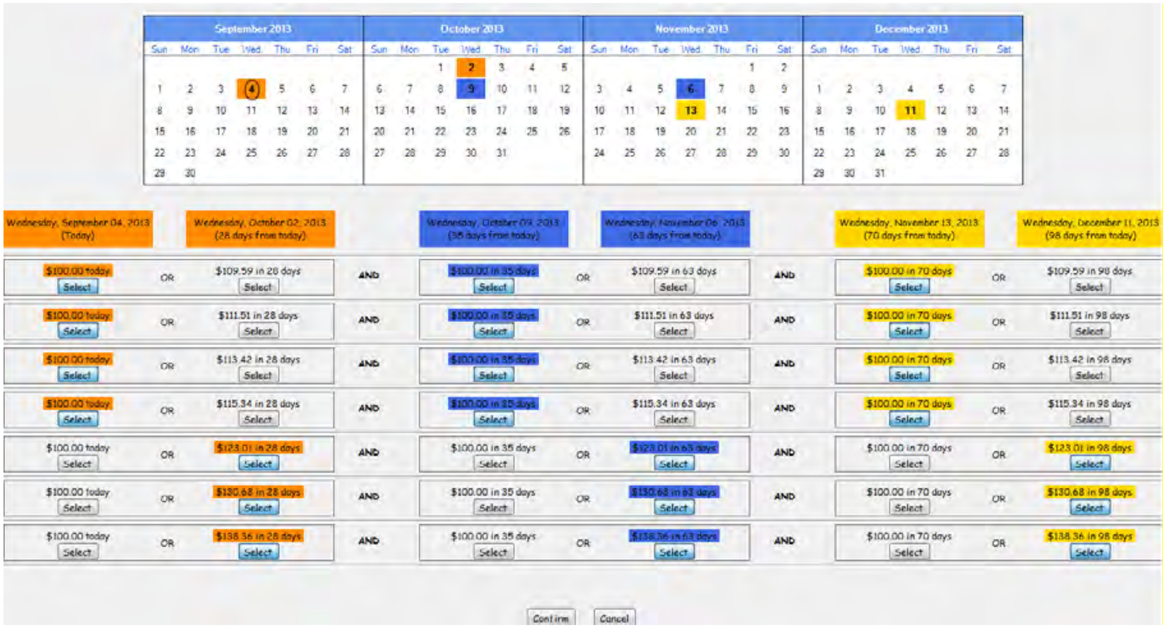

Figure 7: Decision task from treatment T6

The decisions in treatment T6 were *free*. As in treatment T4, subjects could choose different options in a decision row. The instructions emphasized this freedom of choice. Different colors merely indicated different decisions within a triple.

A final treatment T7 implemented a *forced condition* compared to treatment T6, with three decisions in a row. Again, the SS date of the first decision is the day of the session, September, 5th and its LL date is October, 3rd 2013. In the second and third decisions of the triple the two SS dates were October, 10th and November, 14th, whereas the two LL dates were November, 7th and December, 12th 2013, respectively. A typical screen for treatment T7 is presented in Figure 8.

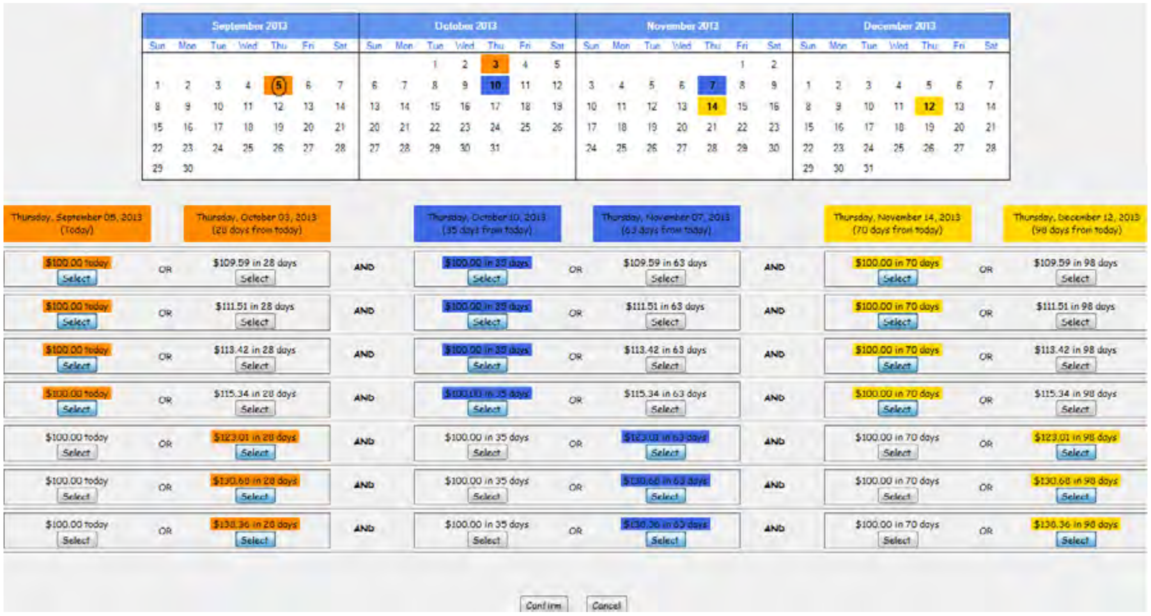

Figure 8: Decision task from treatment T7

Subjects in treatment T7 had to choose the same options in each decision within a single row. The options between rows could differ, as long as they were consistent within one row.

7.2 Recruitment and payment procedures

The subjects were undergraduate students at Georgia State University (GSU). They were recruited by email, after they had previously been addressed in several class sessions across campus where they had been encouraged to sign up and enter the subject pool of the Experimental Economics Center at GSU.[14] The turnout for the eight treatments was as shown in Table 1.

The subjects arrived at the laboratory, without knowing the exact nature of the experimental task. They drew a random number and were seated at a computer station with the

[14] In detail, the recruitment procedure works as follows: a random sample is draw from the students, who have registered in the web-based recruitment system. Then, an invitation to participate in the experiment is sent via email to each student in the random sample. When students logged into the system and confirmed their availability they were selected, until all available seats for the treatment in question were assigned.

corresponding number. Then, detailed instructions were handed to them to review before they were read aloud. The subjects were invited to ask questions to clarify the decision task. In practice, very few questions were ever asked. After the subjects had any questions answered the choice task began. Once the task was completed, the subjects were prompted to answer a demographic survey. The payout selection procedure concluded each experimental session.

A show-up fee of $5 was paid to all subjects, regardless of their performance in the decision tasks. The payouts from the choice task were selected as follows. After completing the task and the subsequent demographic survey, each subject was approached by a research assistant. The subjects were asked to roll a 10-sided die to select one of the ten decision *screens* from the task. After the screen was selected, the subjects were asked to roll the 10-sided die again, until a number between 1 and 7 came up, to select one of the seven decision *rows* on that screen. Their preferred option in the selected row was considered for payout.

To limit experimental costs an additional stochastic procedure determined whether high amounts would be paid out. In treatments T0, T1, and T3, if the preferred option in the selected decision row was less than $100, the amount would be paid for certain. If the selected option was $100 or more, the subjects were asked to roll the 10-sided die one more time. If the outcome was 1, the subjects would be paid this high amount on the date indicated by their choice. If any other number came up, the subject would not be paid for the task. These procedures were explained to subjects in the instructions.

The payment procedure for treatments T4, T5, T6, and T7 was similar. However, if one chosen option was less than $100, the whole decision row (i.e. all preferred options of the selected pair or triple) would be paid to the subject on the chosen dates. If any amount in the selected row was $100 or more, the subjects would have to roll a 10-sided die again. If the outcome was 1, the subjects were paid for all decision made in this row. Otherwise, they would not be paid for the task.

In treatment T2, where the amounts were tripled, a selected amount was considered low if it was less than $300. The payment of high amounts, of $300 or more, was again determined by the same stochastic mechanism.

Only the show-up fee of $5 was paid in cash, at the end of the experimental session. All payouts from the decision tasks were paid using the online payment service *PayPal*. Even if a selected payment was due on the day of the experimental session, the amount was paid via *PayPal* later that day. This common payment procedure is how the transaction costs were kept constant across treatment and choices of SS/LL, so that they would not confound the inferred time preferences.

Whenever first payments were due, the subjects received an email notification one day in advance. No such notification was sent when second or third payments were due, since it was assumed that the subjects had become sufficiently familiar with the procedure. The initiation of each transfer led to another email notification, automatically sent by *PayPal*, which informed the subject that the amount was available and could be claimed. In a few rare cases subjects failed to claim their payments. In these situations, they again received an email notification that explained that they had failed to claim their payments and that *PayPal* would return the fund to the experimenters if the amount remained unclaimed after a period of 30 days. It was pointed out to these subjects that in this case the money would be sent again for a second 30 day period. The *PayPal* transfers would be terminated only if the subjects failed to claim their money a second time. In this case, the subject would have to contact the experimenter, using a previously given email address, to arrange payment. The vast majority of transfers went as planned. Only eight payments were left unclaimed by five different subjects, after the period for all payments had passed. One participant left three payments, another left two payments, and three participants left one payment unclaimed.

8. Results

8.1 Local polynomial regression results

This chapter begins the statistical analysis of the experimental data. This sections presents the results from several local polynomial regression models. The Figures below do usually not contain any confidence intervals (CI) bands. Deviations from general tendencies are more carefully examined in the two following sections. Only when the similarity of choices from different treatment conditions serves as justification for the pooling of choice data is the 95% CI included.

The first local polynomial regression, which concentrates on the choices from treatment T0, is presented in Figure 9. It presents for each magnitude of the principal how the fraction of LL choices varies with the nominal annualized interest rate. The horizontal line at the 50 percent mark indicates where the observed choices are evenly split between SS and LL. Figure 9 illustrates that the fraction of LL choices increases with the nominal annual interest rate. In other words, more people are choosing LL as the interest rate increases. One can also observe the general tendency that the fraction of LL choices increases with the magnitude of the principal. That is, the observed behavior is consistent with a “magnitude effect,” since more subjects are choosing LL when the magnitude of SS is higher. Only for very low and very high annual interest rates, can we see some exceptions from this general tendency. The observed choices in treatment T0 suggest that the descriptive analysis should proceed by examining how the fraction of LL choices varies with nominal annual interest rates for each of the different principal amounts. Moreover, the other Figures in this section will present the local polynomial regression graphs by grouping low principals ($10, $30, and $60) and high principals ($100 and $300) separately.

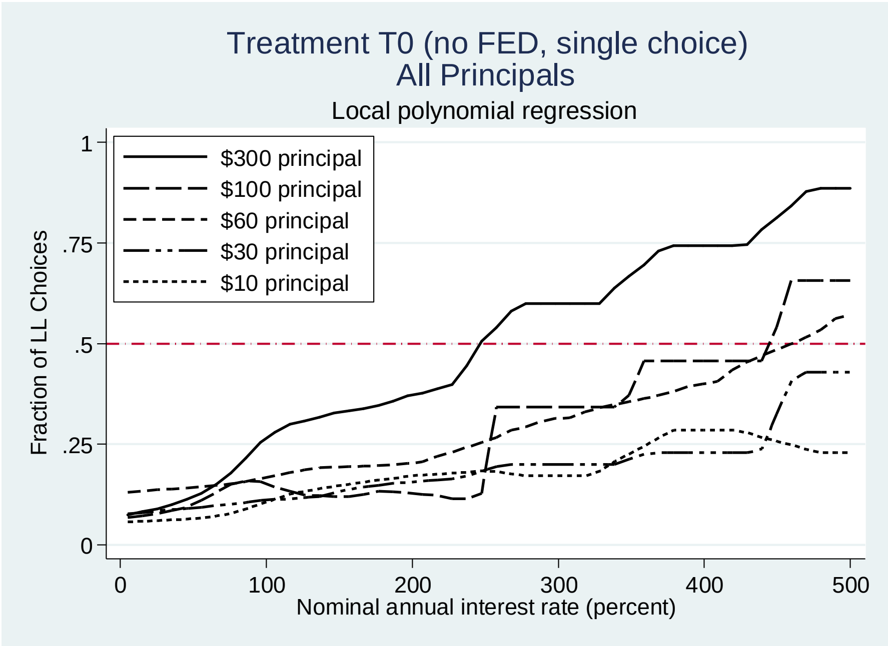

Figure 9: Fraction of LL choices interest rate offered – Treatment T0: All principals

A possible explanation for the apparent "magnitude effect" in Figure 9 is rounding behavior, as discussed in section 3.4.

Such rounding error problems may be particularly pronounced for the lower payoffs, $10 and $30. Moreover, an apparent magnitude effect is also implied by fixed costs discounting, as proposed by Benhabib, Bisin, and Schotter (2010). Andersen, Harrison, Lau, and Rutström (2013, 684) demonstrate that a fixed monetary premium implies not only variable discount rates that decline with the horizon, but also a "magnitude effect." It will not be possible to parse the increase in the fraction of LL choices into the underlying causes before performing a structural analysis.

Before the descriptive analysis attempts to find evidence for the presence of reward bundling, it will be useful to establish that the experimental design managed to ensure that the LL option is chosen at all. If subjects never chose LL in the *control treatments*, T0 to T3, it would not be possible to find any evidence for reward bundling in the *bundling treatments,* T4 to T7. Low annualized interest rates result in relatively low LL rewards. If discount rates are high, the attractiveness of LL rewards is further diminished. In this case, the present utility of even a whole series of LL rewards may not exceed the present utility of a series of SS rewards. Consequently, no reward bundling would be observable. Figures 10a and 10b demonstrate that this is not the case.

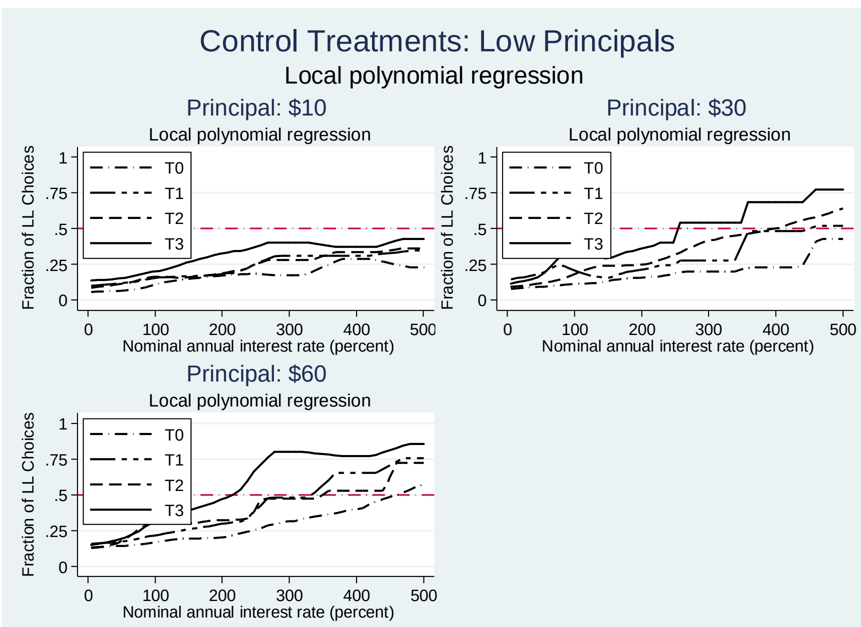

Figure 10a: Fraction of LL choices and interest rate – Control treatments (Low Principals)

Even for low amounts, some patient choice behavior, in which LL is chosen more than 50 percent of the time, is observable in all treatments, although the fraction remains very low with a $10 principal. As the principal increases, one can observe more LL choices for all treatments.

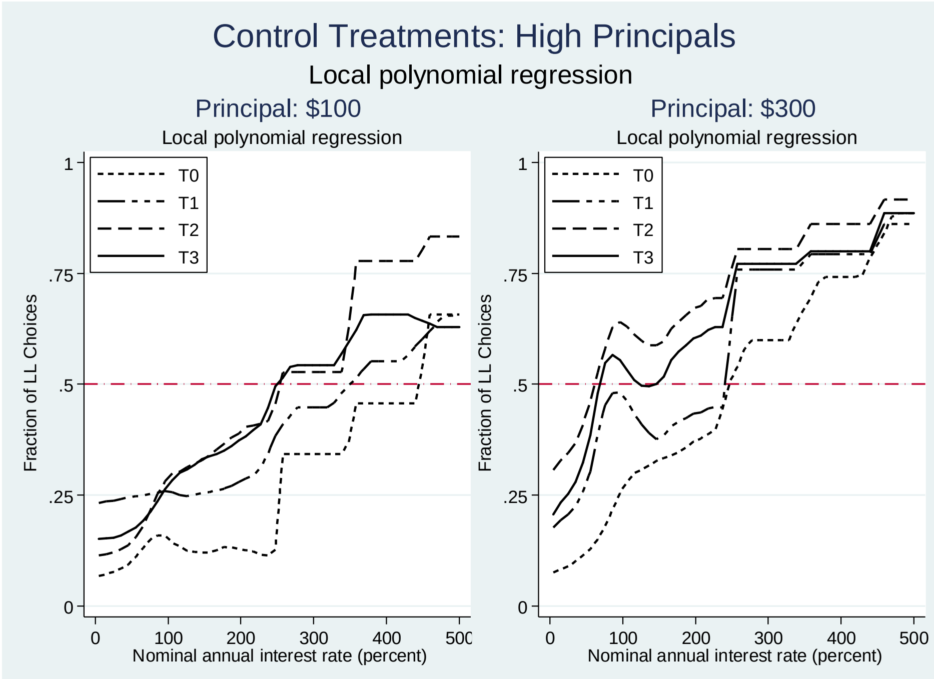

Figure 10b: Fraction of LL choices and interest rate – Control treatments (High Principals)

With the high principals in Figure 10b, even comparatively low interest rates are sufficient to lead to a noteworthy fraction of LL choices in excess of 50 percent.

The choice behavior in the *control treatments* in Figure 10a and 10b also reveals some tentative evidence for a variable discount rate. Subjects in treatment T1, with the 1 day FED, chose more patiently than subjects in T0. A few exceptions are visible with the two high principals in Figure 10b, when the interest rates are very high. The choice behavior in treatment T1 implies a lower discount rate after a FED of 1 day. The descriptive probit analysis in the next section examines whether this possible effect is statistically significant. This behavior is consistent with hyperbolic ($D^{H1}$and D^{H2}) as well as quasi-hyperbolic D^{QH} specifications. However, the visceral interpretation of quasi-hyperbolic discounting predicts that, after the initial drop in utility, the choice behavior will approximate exponential discounting. The observable difference between treatments T1 and T3 is not consistent with this prediction. For low as well as for high principals the fraction of LL choices increases for treatment T3, relative to treatment T1. This suggests that subjects chose more patiently when the FED was increased from 1 day to 35 days. This behavior is more consistent with hyperbolic, rather than quasi-hyperbolic discounting. A more rigorous investigation of the underlying choice behavior will have to be accomplished by a structural estimation of the choice models that were introduced in section 2.1.

The next step is to decide whether there are significant differences between the data from the *free conditions* of the bundling treatments T4 and T6. If no difference can be found data from the four *bundling treatments* (i.e. T4, T5, T6, and T7) can be pooled, when tests for the effects of *bundling treatments* are performed. Figures 11a and 11b focus on the similarity of the first and second choices in treatment T4 and assess whether the pooling of choice data is justified. To aid the assessment of similarity, the 95 percent CI of the first choices is included. Figures 11a and 11b show no significant difference between the first and second choices for low and high principals, respectively.

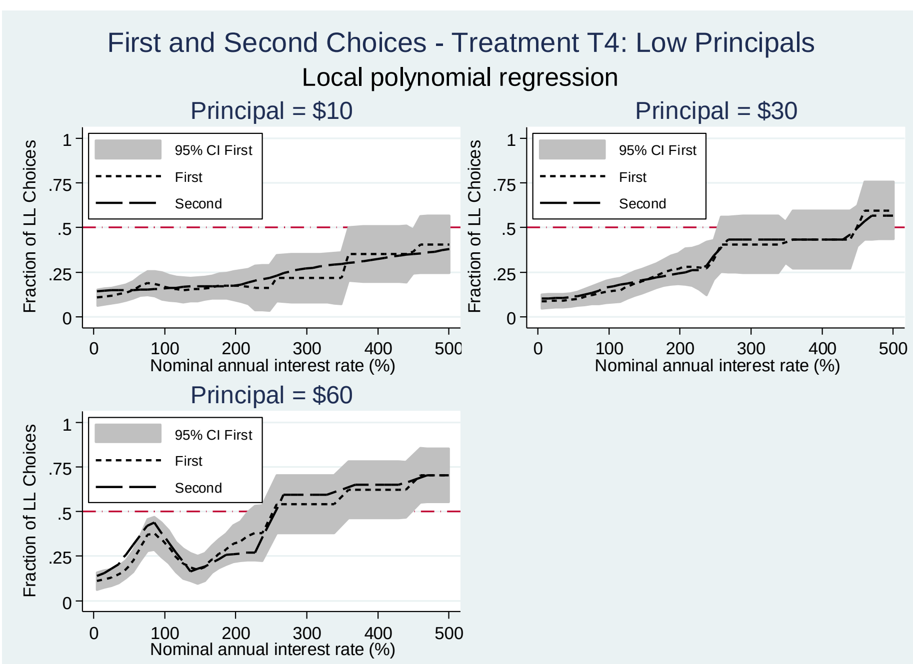

Figure 11a: Comparison between first and second choices in Treatment T4 (Low Principals)

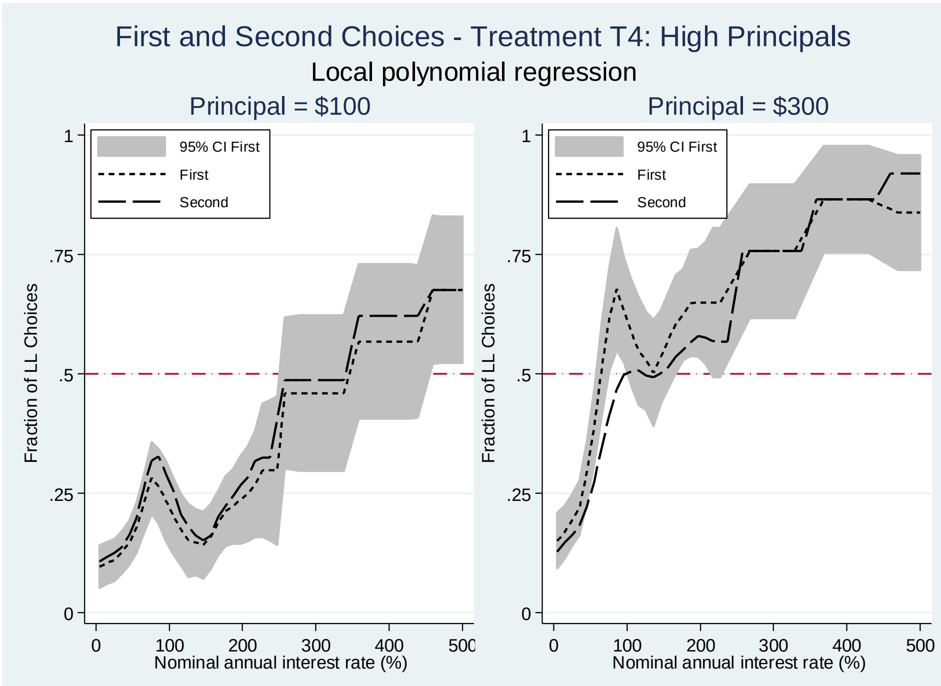

Figure 11b: Comparison between first and second choices in Treatment T4 (High Principals)

The shaded area highlights the 95 percent confidence interval (CI) of the first choices. For all three principal amounts, the second choices lie well within this area interval.

A similar picture emerges for the first, second, and third choices in the other *free condition*, treatment T6. Again to visualize the similarity of choices that are subsequently pooled, the 95 percent CI is added. Figures 12a and 12b suggest the absence of significant differences between first and later choices.

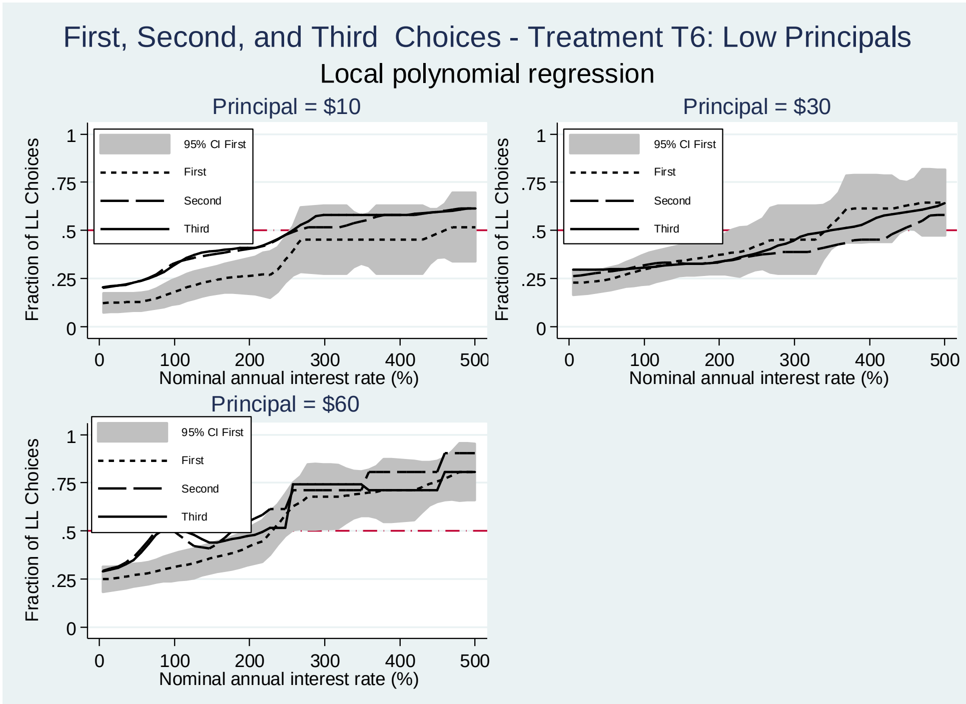

Figure 12a: Comparison between first, second, and third choices in Treatment T6 (Low Principals)

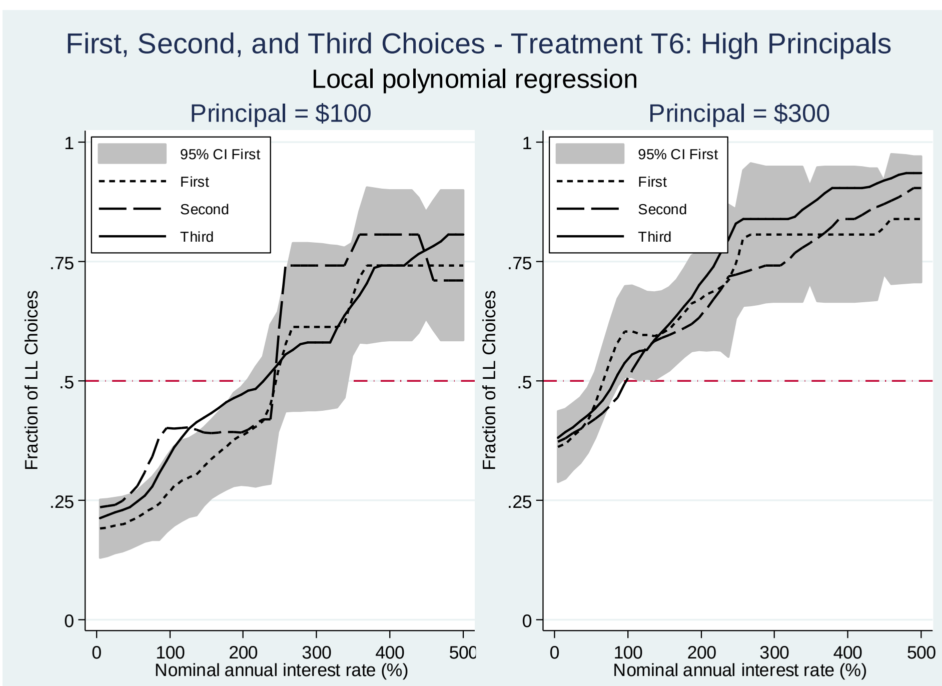

Figure 12b: Comparison between first, second, and third choices in Treatment T6 (High Principals)

Again, the shaded area highlights the 95 percent interval of the first choice. It is apparent from Figures 12a and 12b that the other choice data are mostly within this interval. Slight deviations from this general tendency disappear as the nominal annual interest rate increases. There are no significant differences between the first and later choices in the *free bundling conditions*. Hence, all choices from the *free* treatments T4 and T6 can probably be pooled with choices from the *forced* treatments T5 and T7 when the effects of *bundling treatments* are investigated.

A first glance at the effect of presenting a series of SS/LL decisions is obtained by comparing the choices from the baseline treatment T0 with all *bundling choices* that are pooled across treatments T4, T5, T6, and T7. Since the focus switches back to general tendencies in choice behavior, the CI bands are omitted. Figures 13a and 13b reveal that the observed choice behavior is consistent with the reward bundling hypothesis for low and high payouts, respectively.

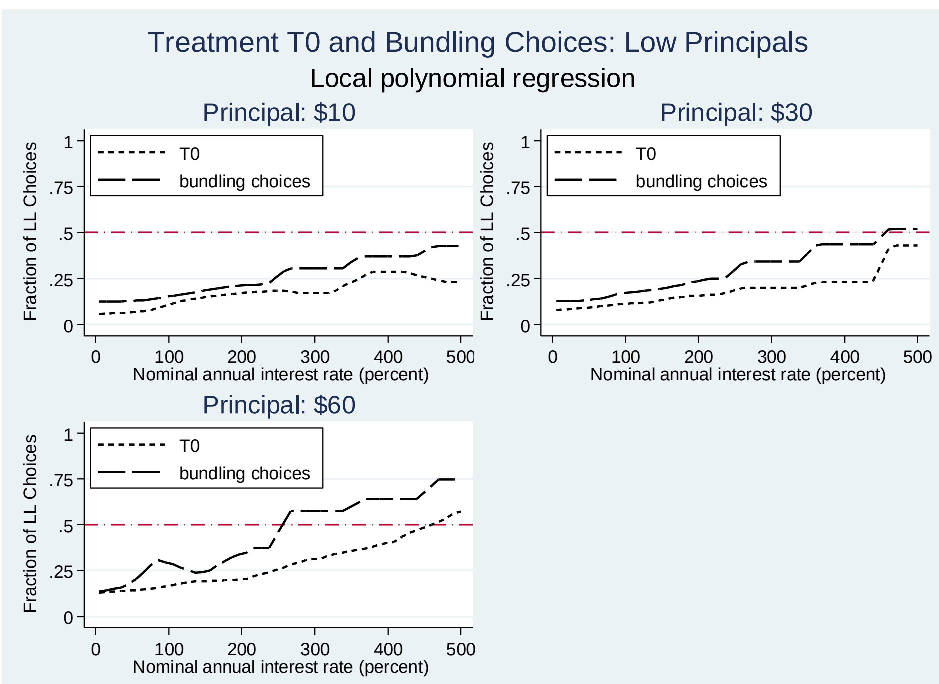

Figure 13a: Comparison of T0 choices and all bundling choices (Low Principals)

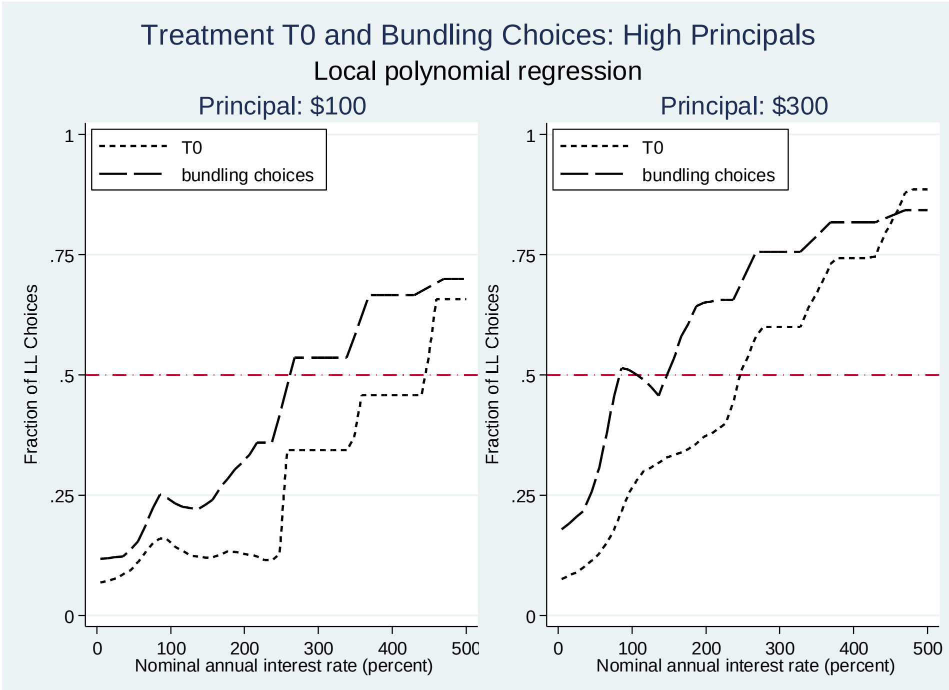

Figure 13b: Comparison of T0 choices and all bundling choices (High Principals)

Choices in treatment T0, which only involves single SS/LL decisions without a FED, exhibit a higher discounting behavior for all interest rates than *bundling choices* in Figure 13a and most interest rates in Figure13b. The only minor exceptions occur for very high interest rates with a $300 principal in the left panel of Figure 13b. This result strongly supports the reward bundling hypothesis. Subjects appear more patient when facing a series of SS/LL decisions, than when being presented with an isolated decision.

Whether reward bundling is present can be further investigated by having a closer look at what choices are affected by the *bundling treatments*. Reward bundling is a mechanism of impulse control. If it is effective, then it should help overcome the temptation of the immediately available SS reward. Therefore, it is informative to see the effect of *bundling treatments* on the *first* choices. Figures 14a and 14b reveal that the fraction of LL choices is indeed higher among the first choices in *bundling treatments* than the comparable choices in treatment T0.

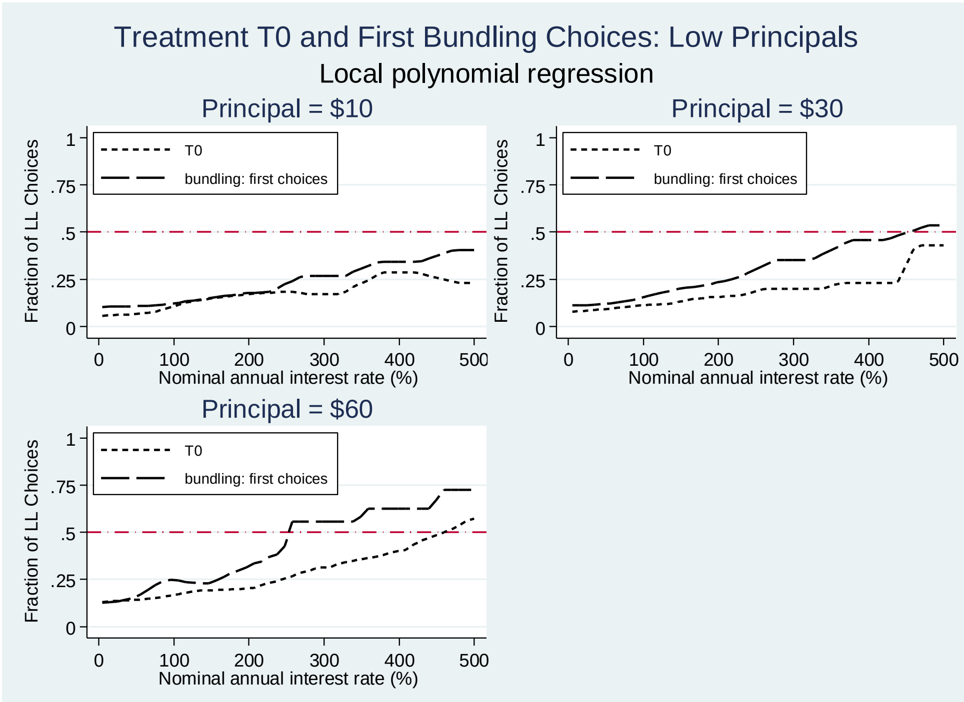

Figure 14a: Comparison of T0 choices and first bundling choices (Low Principals)

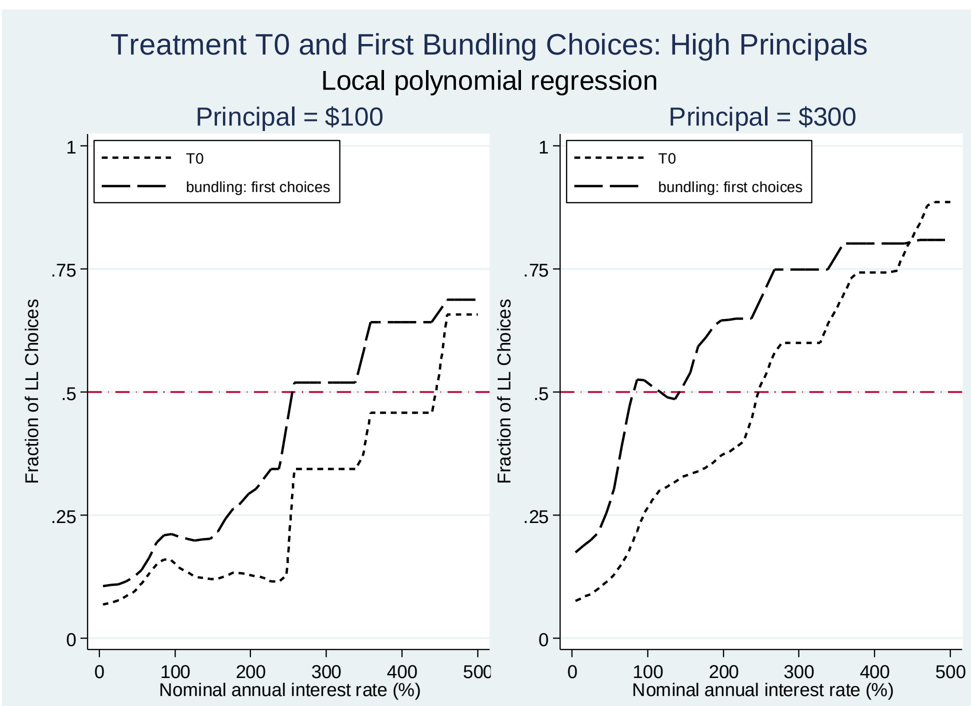

Figure 14b: Comparison of T0 choices and first bundling choices (High Principals)

That first *bundling choices* exhibit higher discount rates than single T0 choices is consistently the case, with all principals. Again, a minor deviation can only be seen in case of the $300 principal, when the interest rates are very high.

Another informative comparison involves the choices with a 35 day FED from treatment T3 and the second choices from the *bundling treatments*. Reward bundling predicts in this case that the choice behavior in treatment T3 should display more patience than second choices from the *bundling treatments*. The reason is that with reward bundling the choices are associated, in treatments T5 and T7 by design and in treatment T4 and T6 by the perceived linkage of the decisions. If subjects bundle second choices, where the reward will be available after a FED of 35 days, with first choices, which do not involve a FED, then this should make choices appear less patient than a situation where they make decision with a 35 day FED in isolation. Consequently, one should expect a higher fraction of LL choice in treatment T3 than in the second choices in the *bundling treatments*. Figures 15a and 15b confirm that this is generally the case.

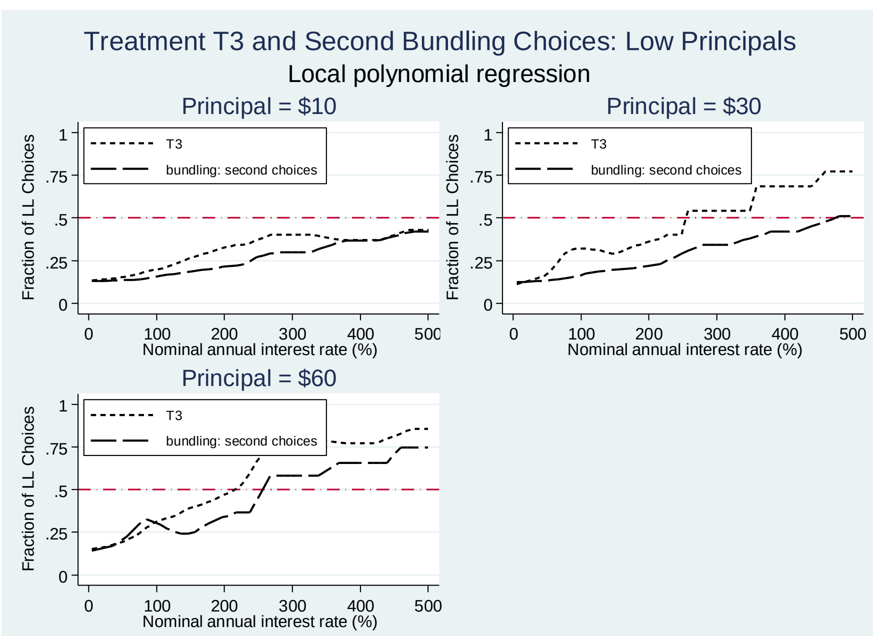

Figure 15a: Comparison of T3 choices and second bundling choices (Low Principals)

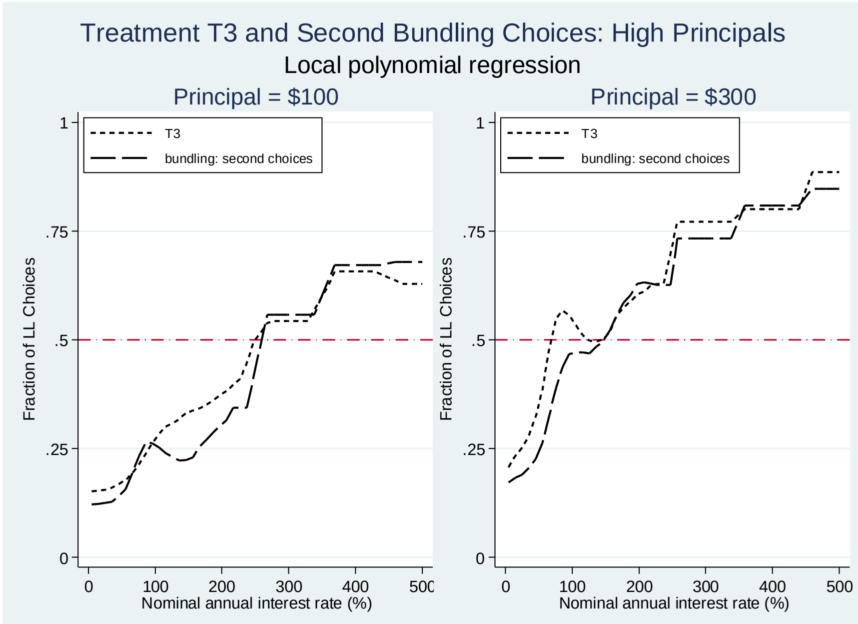

Figure 15b: Comparison of T3 choices and second bundling choices (High Principals)

When the principal amounts are high the comparison of the choices from treatment T3 and the second choices from the *bundling treatments* is less clear. There are some deviations from the predicted behavior that the choices from T3 should display a higher fraction of LL choices. However, in these cases the choices are close together and the discrepancies are not very pronounced.

Another interesting comparison is between the choices in *control treatment* T2, where the rewards are tripled, and the first choices in treatments T6 and T7, which involve three sequential payouts. If subjects bundled rewards, their first choices in those *bundling treatments* should be very similar to the choices in treatment T2.

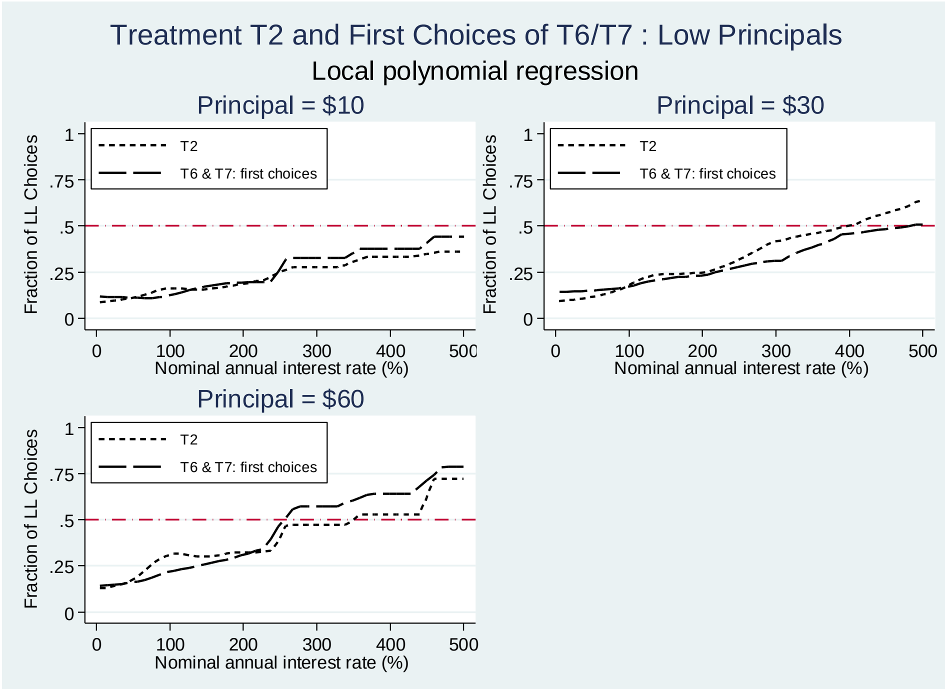

Figure 16a: Comparison of T2 choices and first choices in T6/T7 (Low Principals)

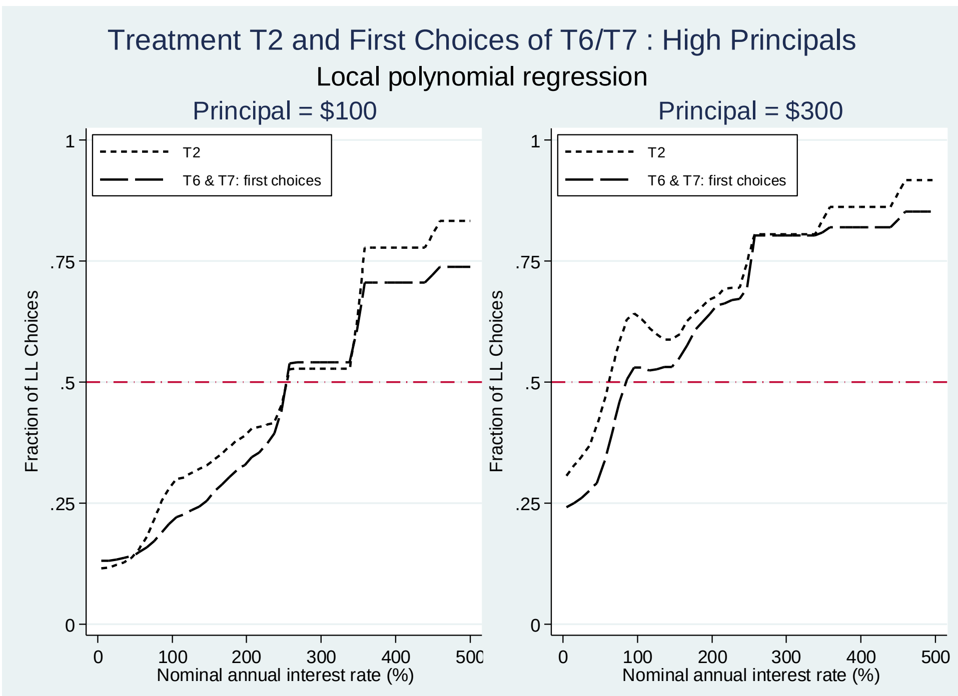

Figure16b: Comparison of T2 choices and first choices in T6/T7 (High Principals)

Figures 16a and 16b show that the choices are indeed very similar which provides additional support for the presence of reward bundling.

The final prediction that is explored in this section is that the fraction of LL choices will increase if the series of SS/LL decisions is extended. The differing effects of bundling on first (Figure 14a and 14b) and second choices (Figure 15a and 15b) suggests a similar distinction when the effects of an additional decision in the series are explored. Figure 17a and 17b compare the first choices in the *bundling treatments* with two decision (T4 and T5) and three decisions (T6 and T7) for low and high principal amounts, respectively.

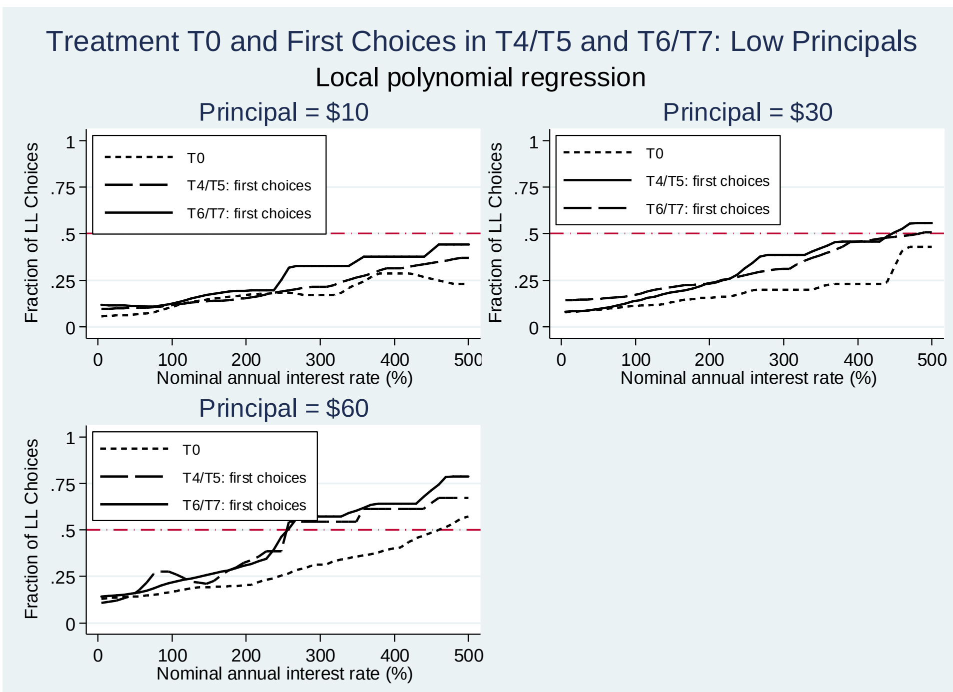

Figure17a: Comparison of treatment T0 choices and the first choices in T4/T5 and T6/T7 (Low Principals)

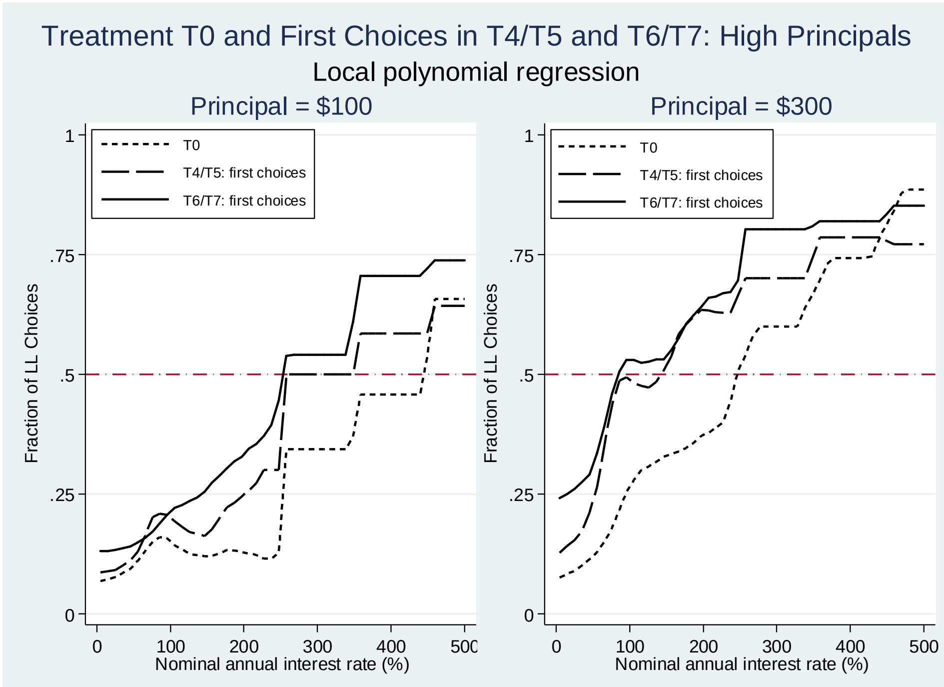

Figure 17b: Comparison of treatment T0 choices and the first choices in T4/T5 and T6/T7 (High Principals)

For most interest rates, there are more LL choices in the *bundling treatments* with three decisions. This is completely consistent with the reward bundling hypothesis. However, there are some minor deviations from this general tendency, especially for low interest rates with low principals in Figure 17a, but also with a $100 principal in Figure 17b. Moreover, when the principals are low, the choices from T4/T5 and T6/T7 are close together, suggesting that the effect of an additional decision in this series is rather weak.

A final comparison involves the second choices from treatments T4/T5 and T6/T7 and the choices from treatment T3. The behavior for low and high principals is presented in Figures 18a and 18b, respectively.

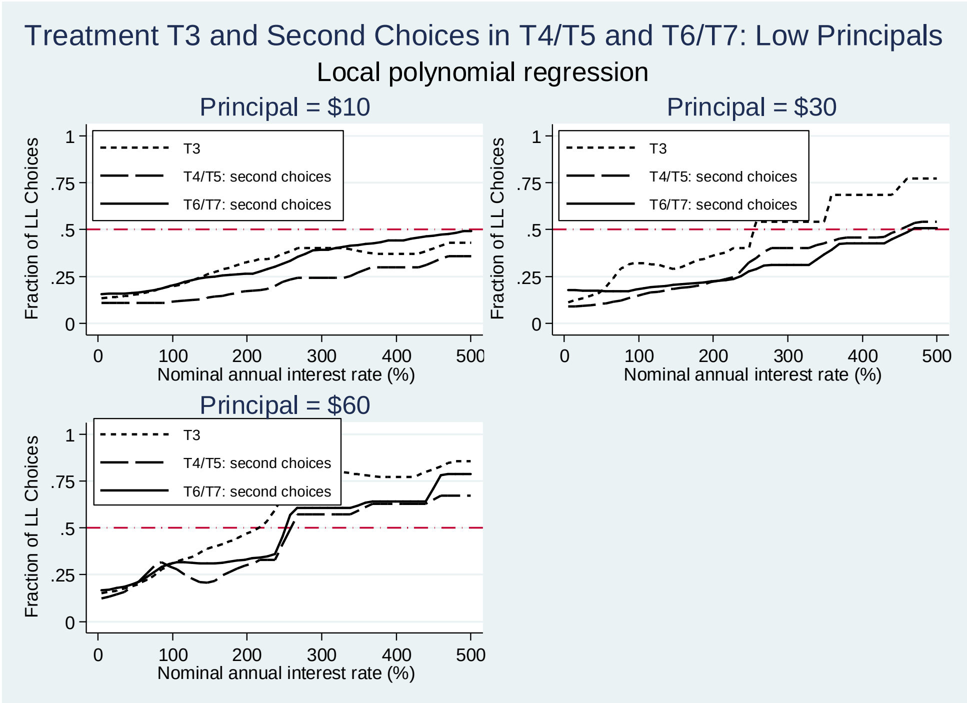

Figure18a: Comparison of treatment T3 choices and the second choices in T4/T5 and T6/T7 (Low Principals)

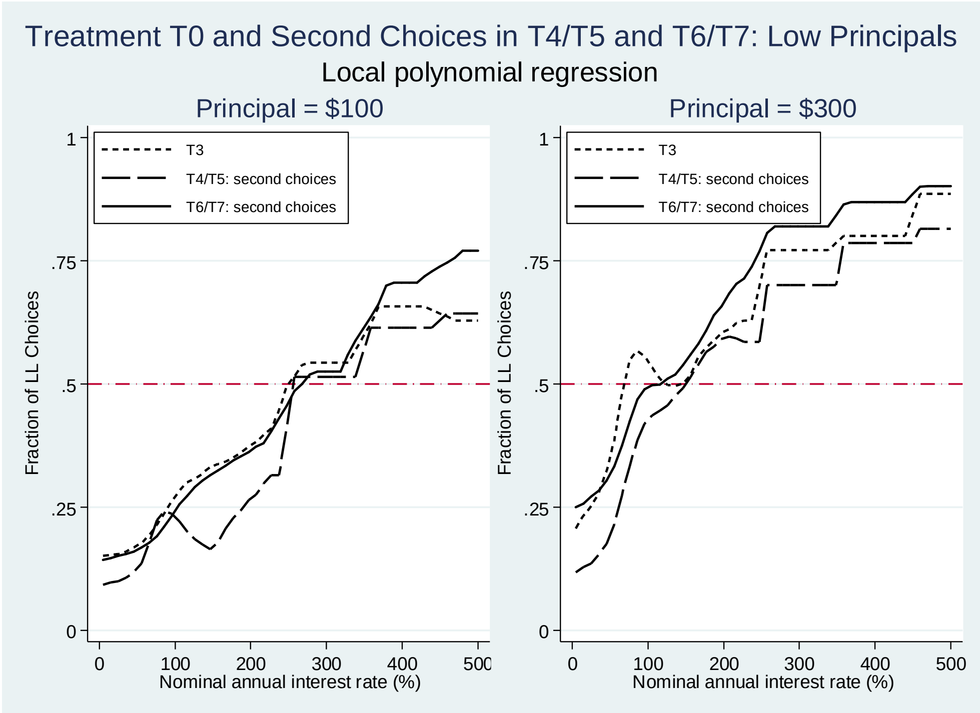

Figure 18b: Comparison of treatment T3 choices and the second choices in T4/T5 and T6/T7 (High Principals)

Except for the low principal of $10, adding a third decision leads to slightly more patient choice behavior. For all other interest rates the fraction of LL choices is generally greater in the *bundling choices* with three treatments. This is consistent with the reward bundling hypothesis, since an additional choice with a longer FED should make the subjects more patient. However, the second choices in both *bundling conditions* should be less patient than in the single choice with a 35 day FED in treatment T3. The reason is that all *bundling treatments* involve making second choices in conjunction with the zero FED condition of the first choices, if reward bundling is applied. This effect is only clearly visible with the low principals of $30 and $60.

Overall, the descriptive explorations in this section provide solid support for the presence of reward bundling. The local polynomial regression analyses reveal tendencies that are clearly in agreement with the predictions stated above, in spite of some slight deviations. However, such broad tendencies in the choice data can only give preliminary evidence that justifies a more thorough investigation. The next section takes the analysis one step further, by running several descriptive probit models.

8.2 Descriptive probit models

In this section several probit models are presented that examine the effects of the treatments on the probability of choosing LL. Like other discrete choice models, probit models make it possible to condition choice probabilities on observable characteristics of the individual decision maker or the choice situation.

Of central importance for the choice situation is the variable BUNDLED, which equals 1 for the *bundling treatments* and 0 for the *control treatments*. Additional binary variables are FIRST, SECOND, and THIRD, which indicate the order of choices in the *bundling treatments*.

The value of the variable FORCED equals 1 if the treatment is T5 or T7 and 0 otherwise. The variable SCALE indicates the effective magnitude of the payoffs. Its value equals 1 in most *control treatments*. Only in treatment T2, where the payoffs are tripled, does SCALE equal 3. In the *forced treatment condition* T5, SCALE takes a value of 2 since every choice effectively leads to twice the chosen payout. Similarly, SCALE equals 3 in the *forced treatment* T7. The variable FED indicates the days of front-end delay that were in effect when each choice was made. Its value equals 0 in treatment T0, 1 in treatment T1, 35 in treatment T3 and the second decisions in the *bundling treatments*, and 70 in the third decision of the *bundling treatments* T6 and T7.

A series of standard socio-demographic questions concluded each of the eight treatment sessions. Among the collected variables is FEMALE, a binary variable that indicates gender. It takes a value of 1 if the subject was female and a value of 0 of the subject was male. The variable BLACK takes a value of 1, if the subject was African American or identified himself as "African" and a value of 0 otherwise. Similarly, the variable ASIAN indicates whether the subject identified his or her race as Asian. Not only the subjects' gender and ethnic background were recorded, their academic performance and ability were also considered relevant. The variable GPAHI indicates a subject that has a grade point average between 3.75 and 4.0, i.e. a "good student," who mostly received A's for letter-graded academic work. Training in the quantitative analysis of decision making was approximated by the variable BUSINESS, which indicated whether the subjects majored in business administration, accounting, finance, or economics.

A first set of probit models analyzes the choice probabilities in the *control treatments*. Of particular interest are the effects of two variables that vary across all control treatments, FED and SCALE. The first marginal effect concerns the front-end delay variable FED. Recall that treatments T0, T1, and T3, introduced FEDs of 0 days, 1 day, and 35 days, respectively. The hyperbolic ($D^{H1}$and D^{H2}) and the quasi-hyperbolic(D^{QH}) specifications predict that the probability of choosing LL increases with the FED. Unlike visceral interpretation of quasi-hyperbolic discounting, the hyperbolic specifications predict no significant difference in the probability of choosing LL between treatments T0 and T1. Hyperbolic specifications also predict that the probability of choosing LL should be higher in treatment T1 than in treatment T3. The polynomial regressions in Figures 10a and 10b, involving all the control treatments, suggest that FED interacts with the magnitude of the principal. Hence, the first probit model includes this interaction term as independent variable. In addition, the probit model also includes interactions of FED with the annualized nominal interest rate, and the monetary premium, i.e. the difference between SS and LL amounts. *Stata* allows the convenient creation of interaction terms by using *factor variables*. Factor variables are extensions of existing variables, which create indicator variables for each value. The remaining independent variables consist of the set of socio-demographic covariates mentioned above.

The average marginal effects (AME) of the 1 and the 35 day FED in treatments T1 and T3, compared to the 0 day FED in treatment T0, are displayed in Figure 19.

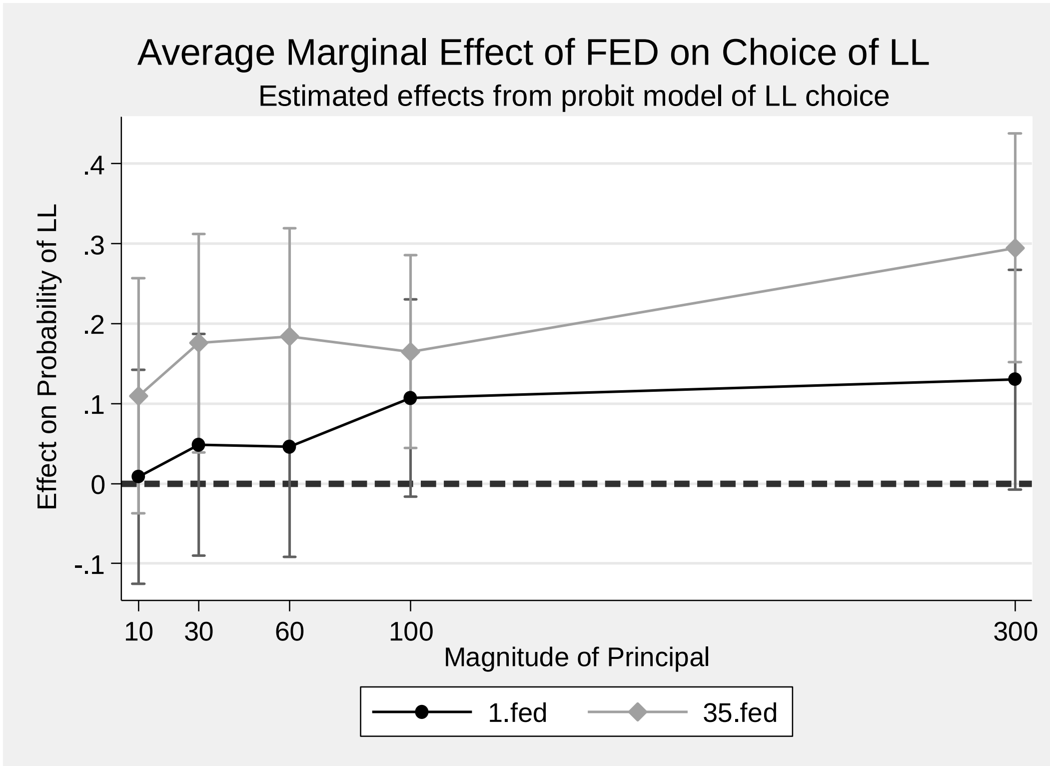

Figure19: Average Marginal Effects of a 1-day FED and a 35-day FED at all Magnitudes of the Principal

One can see a positive AME of the 35 day FED, displayed by the grey line in Figure 19, on probability of choosing LL at all magnitude values of the principal. This positive AME is most pronounced for the highest principal of $300. The 95 percent confidence intervals (CI) of the AME lie mostly above the dashed line at the value of zero that indicates the absence of any AME. One exception occurs with the lowest principal of $10. The AME of the 1-day FED, displayed by the black line in Figure 9, is also above the zero-effect line for all principals. However, the 95 percent CI of the 1-day FED lies partly below zero, even for the highest principal of $300. The values in Table 4 confirm the visual impression.

Table 4: AME of FED at all Magnitudes of the Principal

Variable	AME	Standard Error	*p*-value	95% Confidence Interval	
1 day FED interacted with Magnitude of Principal					
$10	0.009	0.068	0.898	-0.125	0.143
$30	0.049	0.071	0.494	-0.090	0.187
$60	0.046	0.070	0.511	-0.092	0.184
$100	0.107	0.063	0.088	-0.016	0.231
$300	0.130	0.070	0.064	-0.007	0.268
35 day FED interacted with Magnitude of Principal					
$10	0.110	0.075	0.144	-0.037	0.257
$30	0.176	0.070	0.012	0.039	0.312
$60	0.184	0.069	0.008	0.049	0.319
$100	0.165	0.062	0.007	0.044	0.286
$300	0.295	0.073	0.000	0.152	0.438

Panel A of Table 4 reveals that all AME estimates for a 1-day FED are positive. Their values grow with the magnitude of the principal and reach a sizeable point estimate of approximately 10.7 and 13 percentage points for the largest principals of $100 and $300, respectively. However, as the *p*-values indicate, most AME estimates are insignificant at conventional significance levels. Only at the 10 percent level are the AME of a 1-day FED significant, if the principals are $100 or $300. The lack of a significant difference between a FED of 0 days and a FED of 1 day suggests that there is no solid support for the visceral interpretation of quasi-hyperbolic discounting. An alternative quasi-hyperbolic specification, proposed by McClure et.al. (2007), does not restore the plausibility of visceral arousal as the source of impulsive behavior. An additional parameter τmay be able to describe individually varying impatience as a function of delay. However, visceral arousal cannot be a credible explanation for impulsivity after a FED of 1 day, since such arousal can only be maintained for a short time.

Panel B of Table 4 exhibits considerable values for AME point estimates for a FED of 35 days, ranging from approximately 11 percentage points at the lowest principal of $10 to approximately 30 percentage points at the highest principal of $ 300. The *p*-values suggest that the AME estimate of the 35-day FED is not statistically significant at the 10 percent level of significance with the lowest principal of $10. With a principal of $30 the AME is statistically significant at the 5 percent level. With all other principal magnitudes, the AME estimates of the 35-day FED are statistically significant at all conventional significance levels. These results indicate that subjects' inferred discount rate decreases with a 35-day FED in treatment T3, compared to a 0-day FED in treatment T0. Such choice behavior suggests the presence of hyperbolic discounting, since subjects seem to become more patient when the reward is not imminent. The results also provide further evidence against the visceral interpretation of quasi-hyperbolic specification, since the AME of a FED of 35 days differs significantly from the AME with no FED or a FED of 1 day. That people become more patient the more their payouts are delayed is not consistent with the hypothesis of visceral arousal as the sole source of impatience.

The principals in treatment T2 are three times the principals in treatment T0. The variable SCALE indicates the effective magnitude of the payouts. Its value equals 1 in all *control treatments*, with the exception of treatment T2. In treatment T2, the value of SCALE equals 3. This setup allows testing for the presence of an apparent "magnitude effect," by assessing the average marginal effect (AME) of a SCALE of 3. An AME of SCALE=3 on the probability of choosing LL that differs in treatment T2, compared to treatment T0, would be consistent with the presence of a magnitude effect. To investigate whether this is the case is the goal of the second descriptive probit model. As in the first model, the outcome variable is the choice of SS or LL in treatments T0 and T2. The independent variables include the aforementioned demographic variables as well as interactions of SCALE with the annualized nominal interest rate, the monetary premium, and the magnitude of the principal.

The average marginal effects (AME) of SCALE with a value of 3 at all principal magnitudes, compared to SCALE with a value of 1, is displayed in Figure 20.

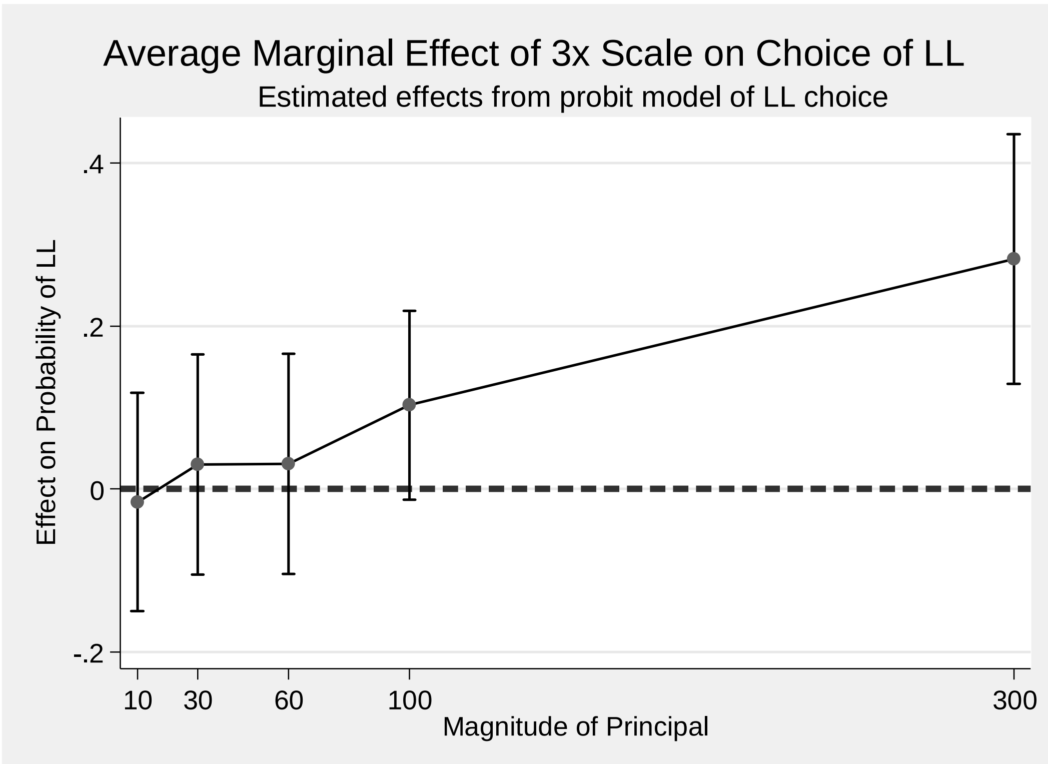

Figure 20: Average Marginal Effects of SCALE=3 at all Magnitudes of the Principal

The point estimate of the AME of a SCALE of 3 is positive at all principals, except for the lowest of $10. However, only in the case of the highest principal of $300 is the 95 percent CI completely above the dashed line in Figure 20, which indicates an AME of 0. This suggests that a significant magnitude effect can only be found for choices that involve the $300 principal. This is further demonstrated by the values in Table 5.

Table 5: AME of SCALE=3 at all Magnitudes of the Principal

Variable	AME	Standard Error	*p*-value	95% Confidence Interval	
SCALE=3 at a Principal of					
\$10	-0.016	0.068	0.812	-0.150	0.118
\$30	0.030	0.069	0.663	-0.105	0.165
\$60	0.031	0.069	0.655	-0.104	0.166
\$100	0.103	0.059	0.082	-0.013	0.219
\$300	0.282	0.078	<0.001	0.129	0.435

The *p*-values in Table 5 reveal that only at a principal of \$300 is the AME of SCALE=3 statistically significant at all customary significance levels. The AME is also statistically significant at a 10 percent significance level, when the principal magnitude is \$100. The polynomial regression in Figure 9 suggested the presence of the apparent “magnitude effect” for lower principal magnitudes. However, the results presented in Table 5 reveal that those are not statistically significant. This finding casts some doubt on the hypothesis that the observed effect is due to rounding behavior. The discussion in subsection 2.2.4 proposes that what appears to be a genuine response to the magnitude of rewards, may actually be the result of rounding to facilitate the comparison of similar amounts. However, the only statistically significant effect occurs with the highest principals, which do not produce similar LL amounts when the annualized interest rates are applied. A final discussion whether the observed behavior constitutes a genuine magnitude effect or is more plausibly explained by rounding is postponed until the structural analysis in the next section.

The analysis can now proceed to the *bundling treatments*. A first tentative test for the presence of reward bundling is accomplished by a simple probit model that focuses on the AME of BUNDLED, which indicates all *bundling treatments*. This model includes the interaction terms of BUNDLED with the magnitude of the principal, the annualized interest rate, and the monetary premium. It also controls for other factors that affect the probability of choosing LL. The preceding analysis revealed that under certain circumstances SCALE and FED can have this effect. They are, therefore, included. Other covariates are SECOND, THIRD, FORCED, and the demographic characteristics mentioned above. Figure 21 reveals the resulting AME of BUNDLED at all principals.

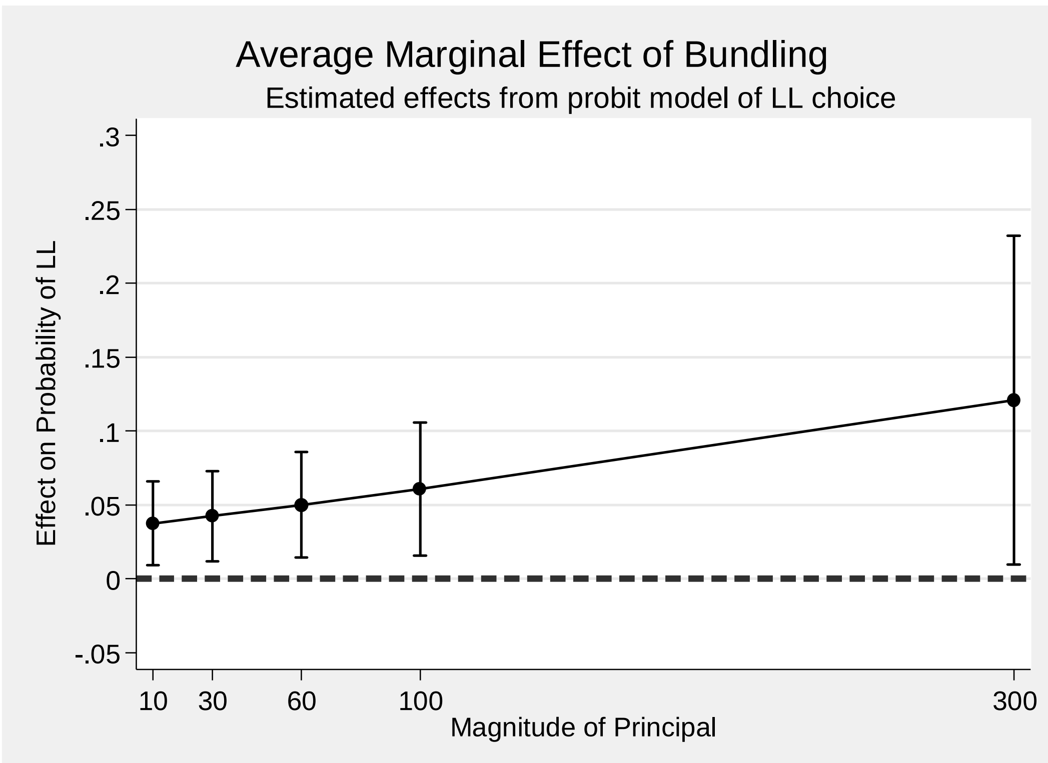

Figure 21: Average Marginal Effects of BUNDLED at all Magnitudes of the Principal

Figure 21 illustrates a positive AME of BUNDLED at all principals, which increases with the magnitude of the principal. One can see that the 95 percent CI is above the dashed line that indicates an AME of zero. The exact values are presented in Table 6.

Table 6: AME of BUNDLED at all Principals

Variable	AME	Standard Error	*p*-value	95% Confidence Interval	
BUNDLED at a Principal of					
$10	0.037	0.014	0.009	0.009	0.066
$30	0.042	0.016	0.007	0.012	0.073
$60	0.050	0.018	0.006	0.014	0.086
$100	0.061	0.023	0.008	0.016	0.106
$300	0.121	0.057	0.033	0.010	0.232

Table 6 confirms that the AME of BUNDLED is significant at all conventional significance levels. It ranges from a rather modest value of approximately 4 percentage points at a principal of $10 to a sizeable effect of 12 percentage points at a principal of $300. These results can be interpreted as tentative support for reward bundling.

The polynomial regression presented in Figures 14a and 14b focused on the first choices in the *bundling treatments*. Likewise, it is interesting to examine how BUNDLED affects the probability of the first choices in *bundling treatments* being LL. For this purpose, another probit model is estimated. It includes interactions of BUNDLED with the magnitude of the principal, the annualized interest rate, and the monetary premium. Furthermore, the model includes the variables SCALE, signifying the effective payout magnitude and FORCED, whether the choices were made in treatment T5 or T7. The model focuses only on the first choices. The resulting AME is presented in Figure 22.

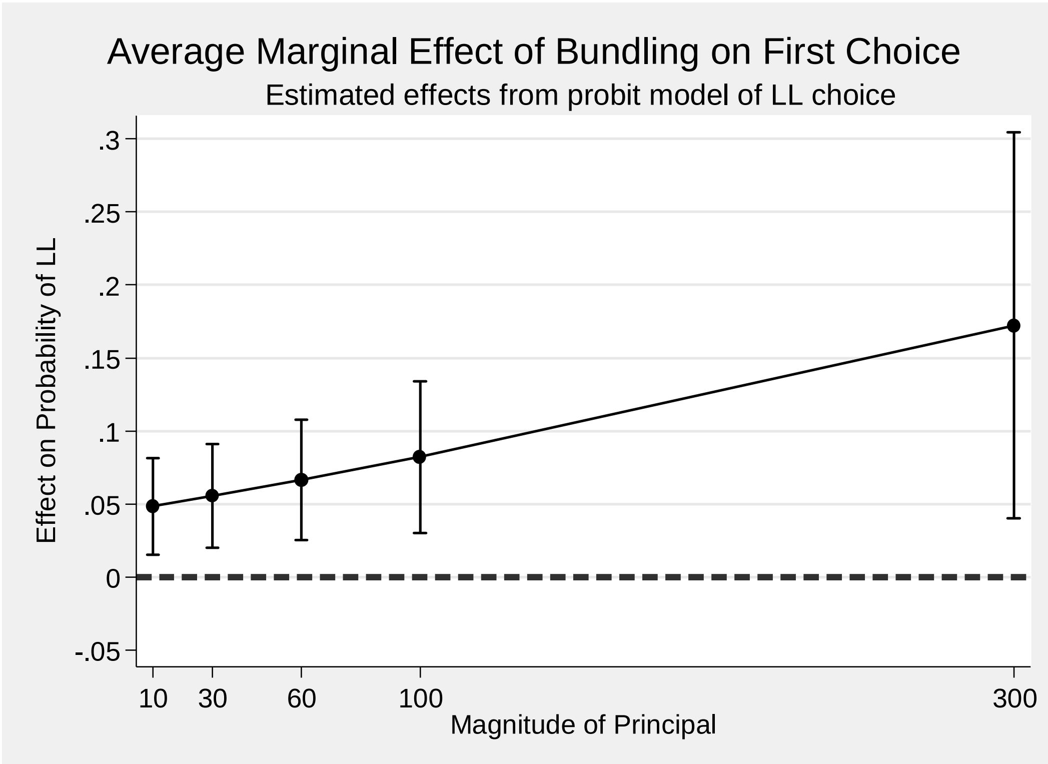

Figure 22: Average Marginal Effects of BUNDLED on first Choices at all Magnitudes of the Principal

Figure 22 shows that BUNDLED has a positive AME on first *bundling* choices at all magnitudes of the principal. The AME increases with the magnitude and the 95 percent CI lies, in all cases, well above the dashed line that indicates an AME of zero. Table 7 presents the exact values.

Table 7: AME of BUNDLED on first Choices at all Principals

Variable	**AME**	**Standard Error**	***p*-value**	**95% Confidence Interval**	
BUNDLED at a Principal of					
$10	0.048	0.017	0.004	0.015	0.081
$30	0.055	0.018	0.002	0.020	0.091
$60	0.067	0.021	0.002	0.025	0.108
$100	0.082	0.027	0.002	0.030	0.134
$300	0.172	0.067	0.011	0.040	0.304

The AME range from 5 to 17 percentage points at the lowest and the highest principals, respectively. A look at the *p*-values in Table 7 reveals that all AME estimates are statistically significant at all conventional levels of significance. These results are consistent with the reward bundling hypothesis.

The descriptive probit models in this section offered further evidence for the reward bundling hypothesis. The next section proceeds to the most advanced level of the econometric investigation and present a comprehensive structural analysis of the experimental data. The structural analysis constitutes the most rigorous test for the presence of reward bundling.

8.3 Structural estimation results

This section presents the results from a comprehensive structural estimation of the experimental data. Unlike descriptive probit models, structural models address how the underlying processes determine the observed choices. Various theoretical models of intertemporal choice were presented in section 2.1. If the prevalence of particular discounting processes differs across treatment conditions, one can infer whether specific treatments facilitate impulse control. In particular, an increased probability of exponential discounting in the *bundling treatments* supports the central hypothesis of reward bundling. Therefore, it is crucial to assess which model is supported by observable choice behavior. If more than one process is at work in a specific treatment, an assumption that seems plausible, it is possible to estimate the relative contribution of each process using a *mixture model*.

As discussed in section 3.1 and chapter 6, discounting parameters must be estimated, conditional on risk parameters, since errors from the risk estimates propagate into the discounting estimates. Consequently, one must collect both, risk and time preference data. The risk preference data for this experiment were not collected by performing a separate task in the same experimental session. Instead, they had been previously collected from samples drawn from the population of GSU undergraduate students. Since the subjects are drawn from the same population one may assume that they are comparable and condition the discounting parameter estimates from one sample on the risk coefficient estimates from the other sample.[15] This procedure is also used by Coller, Harrison, and Rutström (2012, 383-384).The experimental session, producing the risk attitude data, took place in July 2013 and involved a set of binary lottery decisions. The subjects were 171 GSU students, whose recruitment was equivalent to the procedure described in section 7.2. This experiment took place in six sessions over a period of three days.

The structural estimation of risk preference data is discussed in chapter 6. Harrison and Rutström (2008, Appendix F) outline how such models can be estimated using the syntax of *Stata*. Assuming the CRRA specification from (18), this approach is applied to the GSU risk aversion data. Demographic variables are ignored for the moment. The estimation constrained the parameter r to take a value between 0 and 1 and constrained μ to take a positive value. The results are presented in Table 8.

[15] When the likelihood of the discounting parameter is evaluated, conditional on risk parameters, nothing informs the statistical package that the choices over risky lotteries and the choices over SS/LL pairs were made by the same person. One can establish a connection in standard errors, by clustering by the subjects' ID number. However, the clustering over risky choices does not have to match the clustering over time-dated options. So, these choices do not have to come from the same sample, although they should ideally come from the same population (Coller, Harrison, & Rutström 2012, 383-384).

Table 8: Structural Estimation of Risk Data

Parameter	Point Estimate	Standard Error	*p*-value	95% Confidence Interval	
A . EUT					
Log-Likelihood	-8340.471				
r	0.826	0.010	<0.001	0.806	0.846
μ	0.141	0.006	<0.001	0.129	0.153
B. RDU with power function					
Log-Likelihood	-8329.190				
r	0.814	0.011	<0.001	0.792	0.835
γ	1.143	0.075	<0.001	0.996	1.289
μ	0.147	0.007	<0.001	0.133	0.160
C. RDU with inverse-S function					
Log-Likelihood	-8167.272				
r	0.683	0.014	<0.001	0.655	0.710
γ	0.719	0.019	<0.001	0.683	0.756
μ	0.141	0.007	<0.001	0.129	0.154
D. RDU with Prelec function					
Log-Likelihood	-8142.502				
r	0.662	0.014	<0.001	0.634	0.690
η	0.913	0.030	<0.001	0.854	0.971
ϕ	0.623	0.026	<0.001	0.572	0.675
μ	0.140	0.006	<0.001	0.127	0.152

The ML estimation of an EUT model, with CRRA specification and contextual utility, yields the results presented in panel A of Table 8.Both parameters r and μ are statistically highly significant, in the sense of being different from zero.

The estimation results of an RDU model with a power probability weighting function (pwf) are presented in panel B of Table 8. The parameters were again constrained to yield results that are theoretically meaningful. All estimates are statistically different from zero. The parameter r has an estimated value of 0.81, which is only slightly lower than the estimate from the EUT analysis in panel A. This suggest that the concave curvature of the utility function is almost to an equal extent responsible for the implied risk premium as suggested by EUT. In addition, panel B exhibits a value of 1.14 for the parameter γ. Using the power specification for the weighting function, a value greater than 1 suggests that subjects overweight the probabilities of comparatively “bad” outcomes occurring, which can be described as “probability pessimism.” In this case, both concave utility *and* probability pessimism produce the observable risk averse choice behavior. The concepts of probability weights and the implied decision weights are illustrated in Figure 23.

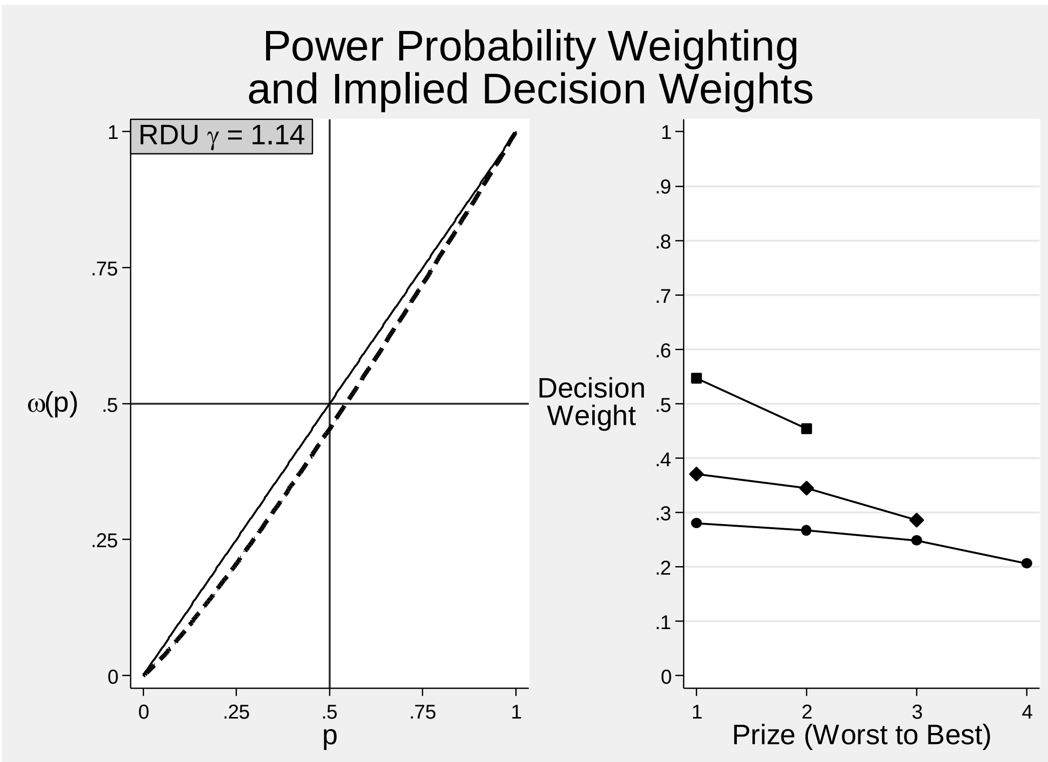

Figure 23: Probability weighting with parameter estimate of the power function and implied decision weights

The left panel displays the probability weighting function, where the objective probabilities on the bottom axis are transformed by $\omega(p) = p^{1.14}$. The straight diagonal line represents the identity function, where objective probabilities p are not distorted, as in EUT. One can see that $\omega(p)$ lies slightly below this 45º line, especially as one moves away from the extremes of 0 and 1. The decision weights implied by this probability weighting behavior are illustrated in the right panel of Table 23. The panel depicts the ranked (worst to best) prizes of reference lotteries, which occur with uniform probability to pedagogically show the pure effect of probability weighting (i.e. when there are 2 prizes, each prize occurs with $p = 1/2$, when there are 3 prizes each prize occurs with probability $p = 1/3$ and so on). Three scenarios are illustrated, with 2, 3, and 4 prizes. In each case, relatively more weight is attached to the worse outcome, than to the subsequent better outcomes. This is a clear case of probability pessimism.

However, one must still examine whether γ is significantly different from 1 in a statistical sense, the case where the specification collapses to EUT. Otherwise, one cannot conclude that probability weighting significantly affects the choice behavior from the risk preference data. A formal hypothesis test reveals that the hypothesis $\gamma = 1$ cannot be rejected. Its p-value falls slightly short of the 0.05 level of significance (*p*-value = 0.056). This result is already suggested by a closer look at the 95 percent CI of γ in panel B of Table 8.
To check the robustness of this result, two additional probability weighting functions are considered. Each is frequently used in applied research.

Panel C of Table 8 presents the result from an RDU estimation of the risk choice data with an inverse-S specification of the probability weighting function. The estimate of $r$is now 0.68, lower than with EUT or RDU with a power specification. The parameter γ is estimated to be 0.72, which does suggest overweighting of small and underweighting of large probabilities. Figure 24 illustrates again the probability and decision weights implied by these estimates.

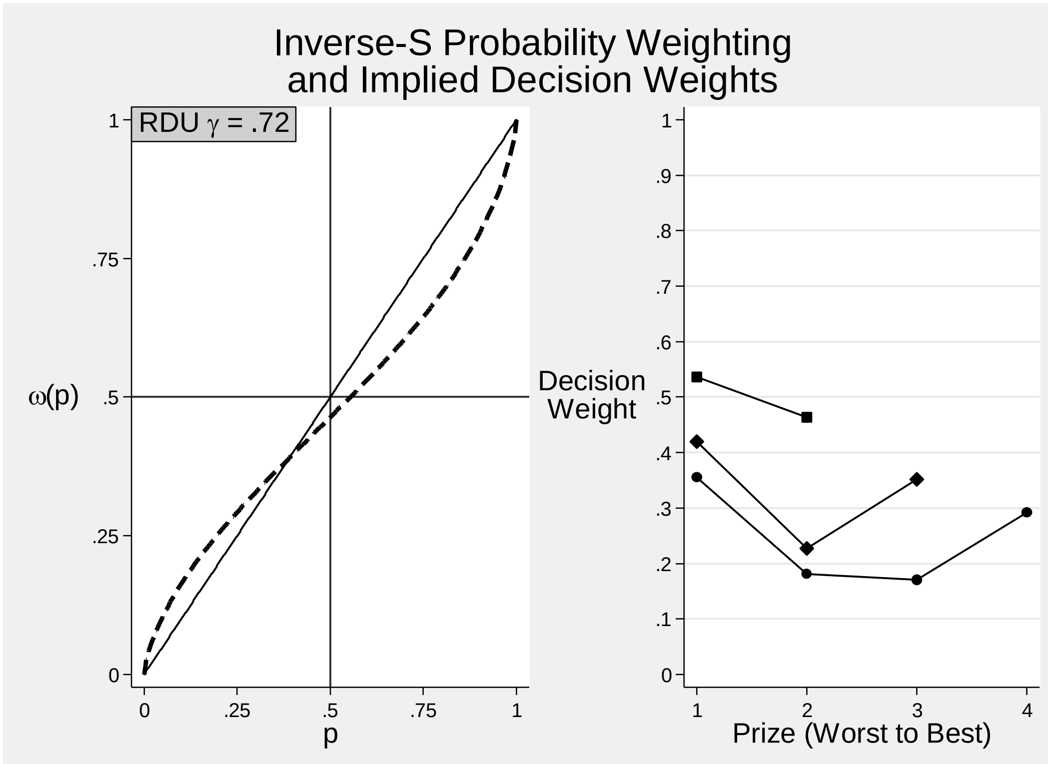

Figure 24: Probability weighting with parameter estimate of the inverse-S function and implied decision weights

The probability weighting function $\omega(p)$ with $\gamma = 0.72$, follows the proposed pattern of overweighting small and underweighting large probabilities. One must then assess whether the estimated value of γ is significantly different from 1. Otherwise, the choice behavior would still be consistent with EUT. However, a formal test shows that the hypothesis $\gamma = 1$ can clearly be rejected (p-value <0.001). Unlike the estimation of the power specification, the estimation of the inverse-S function does indeed produce statistically significant evidence of probability weighting.

The results from the estimation of the Prelec function are presented in panel D of Table 8. The estimate r is 0.66, lower than with the previous specifications. The estimates of η and ϕ equal 0.91 and 0.62, respectively. Figure 25 illustrates the probability weights, resulting from these estimates and the implied decision weights.

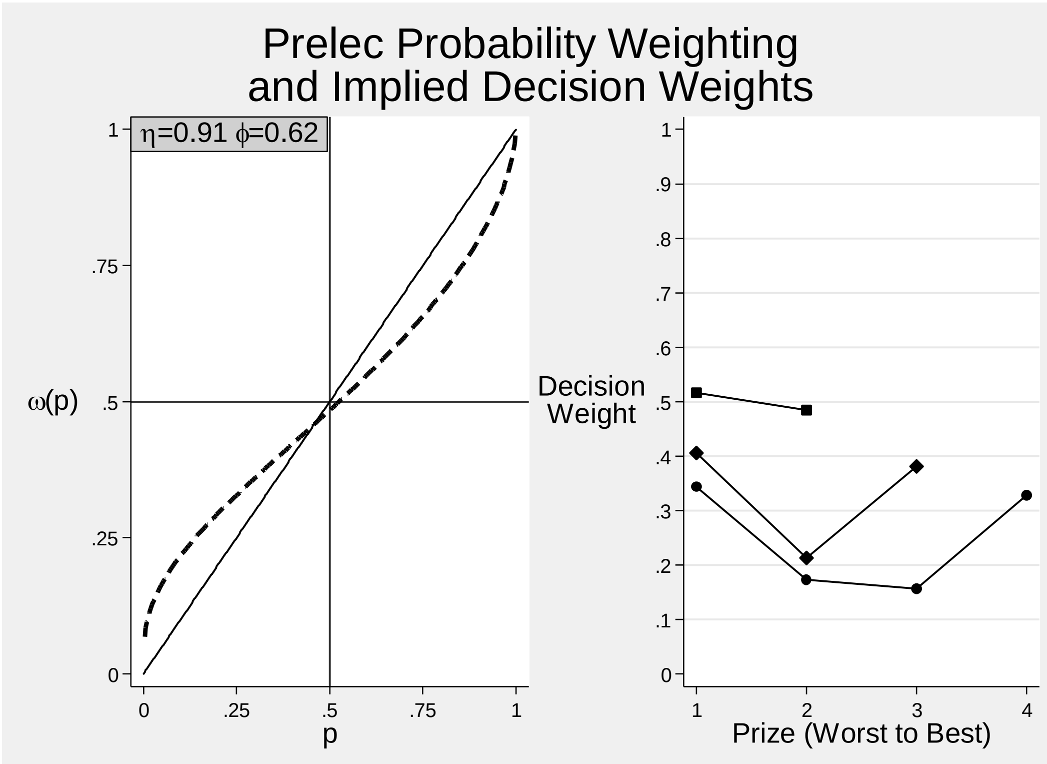

Figure 25: Probability weighting with parameter estimates of the Prelec function and implied decision weights

The patterns of probability and decision weighting resemble those that were implied by the inverse-S specification. It is once again particularly important to assess whether the specification reduces to EUT. A formal test shows that the joint hypothesis $\eta = \phi = 1$ can be rejected with high confidence (*p*-value <0.001). The estimation of the Prelec specification does, therefore, yield additional evidence for probability weighting.

Of all weighting functions estimated, only the power specification failed to reveal any significant probability weighting. The estimation of the inverse-S and of the Prelec functions produced statistically significant evidence for probability weighting. This suggests that some of the implied risk premium can be attributed to probability pessimism, rather than solely attributed to the curvature of the utility function. This becomes obvious when one regards the lower levels of r in the RDU specifications, compared to EUT. Obtaining an accurate estimate of r is of particular importance for the estimation of time preferences, since the curvature of utility function directly influences inferred discount rates. Based on the statistical significance of the estimates of core parameters and the log-likelihood, the RDU specification with Prelec probability weighting appears to be the most convincing of all models considered. Therefore, for the remainder of this chapter it will be the risk attitude model used for the joint estimation of all time preference models examined.

The decision to use previously collected risk attitude data, instead of presenting the same sample of subjects with a risk and a time preference task, has additional implications for the data analysis. The sample of GSU students that were given the risk attitude task produced a considerably smaller number of choice observations than the separate sample that was presented with the eight bundling treatments. The risk preference data consisted of 13,018 choices, 28.88 percent of all data points, whereas the bundling data set contained 32,060 choices, or 71.12 percent of all data points. To give the risk attitude data sufficient influence when the joint estimation of time preferences is performed, they should be weighted. A simulation that evaluated the effect of different weight values was performed. The goal was to find a value that would allow a significant effect on estimated discount rates without drastically altering the estimated risk parameters when the discounting choices are added. Based on these simulation results, a weight of 30 was attached to the risk attitude choices.

The principles of performing structural estimation of time preferences are now demonstrated by analyzing the data from treatment T0. Assume the choices are generated by exponential discounting, with RDU consistent risk preferences with a Prelec probability weighting function. Table 9 presents the results.

Table 9 – Treatment T0: Structural Estimation of the Exponential Discounting Model

Parameter	Point Estimate	Standard Error	*p*-value	95% Confidence Interval	
r	0.666	0.014	<0.001	0.638	0.693
η	0.913	0.030	<0.001	0.854	0.971
φ	0.625	0.026	<0.001	0.574	0.677
δ	2.120	0.335	<0.001	1.464	2.777
μ	0.140	0.006	<0.001	0.127	0.152
ν	0.808	0.116	<0.001	0.580	1.036

All estimates are statistically significantly different from zero at all conventional levels of significance. The value of the parameter r is very similar to the estimate presented in Panel D of Table 8, as one would expect. This suggests that subjects exhibit a significant risk aversion compared to individuals with risk neutral preferences. The parameter δ indicates an annualized discount rate of 212 %.

A slight econometric extension makes it possible to see how key parameters are determined by individual characteristics. Various demographic variables were collected at the end of the experimental sessions. Table 10 illustrates the estimated marginal effects of various demographic characteristics on the parameters δ in treatment T0:

Table 10: Treatment T0 – Marginal Effects of Demographic Covariates on Parameter in the Exponential Discounting Model

Parameter	Variable	Point Estimate	Standard Error	*p*-value	95% Confidence Interval	
r	FEMALE	0.015	0.034	0.660	-0.051	0.081
	BLACK	0.004	0.034	0.910	-0.063	0.071
	ASIAN	0.002	0.057	0.975	-0.110	0.114
	GPAHI	-0.011	0.035	0.762	-0.080	0.058
	BUSINESS	-0.001	0.050	0.988	-0.099	0.097
η	FEMALE	0.099	0.071	0.166	-0.041	0.238
	BLACK	-0.096	0.086	0.259	-0.264	0.071
	ASIAN	-0.127	0.097	0.191	-0.316	0.063
	GPAHI	-0.055	0.065	0.396	-0.182	0.072
	BUSINESS	-0.067	0.077	0.385	-0.219	0.085
ϕ	FEMALE	-0.045	0.049	0.354	-0.141	0.051
	BLACK	0.098	0.067	0.148	-0.035	0.230
	ASIAN	0.166	0.099	0.093	-0.028	0.360
	GPAHI	-0.015	0.054	0.780	-0.120	0.090
	BUSINESS	-0.058	0.069	0.401	-0.193	0.077
δ	FEMALE	0.032	1.019	0.975	-1.966	2.030
	BLACK	0.439	0.967	0.650	-1.456	2.334
	ASIAN	-0.884	0.945	0.350	-2.737	0.969
	GPAHI	-0.390	0.628	0.534	-1.621	0.840
	BUSINESS	-0.048	1.025	0.962	-2.058	1.961

The definitions of the demographic variables, all of which are binary, are presented in section 4.2. No individual marginal effect parameters are significantly different from zero at the 5 percent level. One can test the joint significance of the demographic covariates on the model parameters by performing Wald tests of non-linear restrictions.[16] Three tests are performed to test the joint influence of the demographic covariates on the risk parameters r, η, ϕ, the discounting parameter δ, as well as on the risk and discounting parameters together. The resulting *p*-values are 0.51, 0.66 and 0.61, respectively. Therefore, one cannot conclude that the joint influence of the demographics is significantly different from zero.

[16] Depending how the tests is set up they can be implemented with either the *test* or the *testnl* command in Stata.

Task characteristics, as well as personal covariates, may affect parameter values. Figure 9 suggest that the magnitude of the principal has an important effect on the inferred discount rates. Table 11 presents the marginal effects of different magnitudes of the principal in the discounting tasks. All effects are relative to the highest principal of $300.

Table 11: Treatment T0 – Marginal Effects of the Magnitude of the Principal on the Discount Rate in the Exponential Discounting Model

Parameter	Magnitude of the Principal	Point Estimate	Standard Error	*p*-value	95% Confidence Interval		Discount rate
δ	$300.00 (base category)	1.160	0.206	<0.001	0.755	1.564	116%
	$100.00	1.391	0.369	<0.001	0.667	2.115	255%
	$60.00	1.405	0.500	0.005	0.424	2.385	256%
	$30.00	4.155	1.426	0.004	1.360	6.951	531%
	$10.00	11.588	5.088	0.023	1.616	21.559	1275%

Table 11 shows that the estimated annualized discount rate of the base category, the highest principal of $300, is 116%. This estimate is clearly significantly different from zero (*p*-value <0.001). The point estimates, standard errors, *p*-values, and confidence intervals of the other principal amounts refer to effects on the discount rate δ, relative to the highest principal of $300. The implied discount rates are, of course, constant for all horizons and are reported in the last column. All point estimates of the marginal effects are significantly different from zero. All marginal effects are positive, meaning that discount rates increase, and subjects become more impatient, if the magnitude of the principal is lowered. The marginal effect ranges from approximately 139 percentage points with a principal of $100 to 1159 percentage points with the lowest principal of $10. The resulting annualized discount rates range from 255% to 1275%.

A hyperbolic discounting model is presented in Table 12. Demographic covariates are ignored at this time.

Table 12: Treatment T0 – Structural Estimation of the Hyperbolic Discounting Model

	Point Estimate	Standard Error	*p*-value	95% Confidence Interval	
r	0.666	0.014	<0.001	0.638	0.693
η	0.913	0.030	<0.001	0.854	0.971
ϕ	0.625	0.026	<0.001	0.574	0.677
K	1.189	0.117	<0.001	0.960	1.419
μ	0.140	0.006	<0.001	0.127	0.152
ν	0.808	0.116	<0.001	0.579	1.036

The results in Table 12 show that the estimate of K is 1.19. Hyperbolic discounting implies non-constant discount rates that decline with the horizon. The discount rate that is implied by the D^{H2} specification is derived in (5). Table 13 exhibits this implied discount for several horizons:

Table 13: Treatment T0 – Discount Rates for Several Horizons Implied by Hyperbolic Discounting

Horizon	Point Estimate	Standard Error	*p*-value	95% Confidence Interval	
1 day	2.278	0.382	<0.001	1.528	3.027
2 days	2.272	0.381	<0.001	1.526	3.017
3 days	2.265	0.379	<0.001	1.523	3.007
4 days	2.259	0.377	<0.001	1.521	2.997
5 days	2.253	0.375	<0.001	1.518	2.987
6 days	2.247	0.373	<0.001	1.516	2.977
7 days	2.241	0.371	<0.001	1.514	2.968
2 weeks	2.192	0.356	<0.001	1.494	2.890
1 month	2.107	0.331	<0.001	1.459	2.756
2 months	1.959	0.289	<0.001	1.392	2.526
3 months	1.832	0.256	<0.001	1.331	2.333
6 months	1.543	0.187	<0.001	1.177	1.908
1 year	1.189	0.117	<0.001	0.960	1.419
2 years	0.838	0.064	<0.001	0.713	0.963
5 years	0.473	0.025	<0.001	0.425	0.522

Table 13 shows that the implied annualized discount rate declines from 228% with a horizon of 1 day to 47% when the horizon is extended to 5 years. A notable decrease in the value of the discount rate can be observed from a horizon of 1 month on. Between a horizon of 1 month and a horizon of 6 months, the discount rate drops drastically from 211% to 154%.

One can also allow the key parameters of the hyperbolic discounting model to be determined by demographic characteristics. This is illustrated in Table 14:

Table 14: Treatment T0 – Marginal Effects of Demographic Covariates on Parameters in the Hyperbolic Discounting Model

Parameter	Variable	Point Estimate	Standard Error	*p*-value	95% Confidence Interval	
r	FEMALE	0.014	0.034	0.670	-0.052	0.081
	BLACK	0.004	0.034	0.912	-0.064	0.071
	ASIAN	0.002	0.057	0.965	-0.109	0.114
	GPAHI	-0.010	0.035	0.779	-0.079	0.059
	BUSINESS	-0.001	0.050	0.979	-0.099	0.096
η	FEMALE	0.099	0.071	0.166	-0.041	0.238
	BLACK	-0.097	0.086	0.259	-0.264	0.071
	ASIAN	-0.127	0.097	0.190	-0.316	0.063
	GPAHI	-0.055	0.065	0.396	-0.182	0.072
	BUSINESS	-0.067	0.077	0.385	-0.219	0.085
ϕ	FEMALE	-0.046	0.049	0.352	-0.141	0.050
	BLACK	0.098	0.067	0.148	-0.035	0.230
	ASIAN	0.166	0.099	0.093	-0.028	0.360
	GPAHI	-0.015	0.054	0.784	-0.120	0.090
	BUSINESS	-0.058	0.069	0.399	-0.193	0.077
K	FEMALE	0.087	0.376	0.816	-0.649	0.823
	BLACK	0.123	0.313	0.694	-0.491	0.737
	ASIAN	-0.351	0.367	0.339	-1.069	0.368
	GPAHI	-0.156	0.239	0.513	-0.625	0.312
	BUSINESS	-0.021	0.370	0.956	-0.745	0.704

Again, none of the demographic variables is individually significantly different from zero. This resembles the pattern that emerged when the exponential discounting model was estimated. It is tested whether the demographic variables jointly affect the risk parameters $r, \eta, \phi,$ the discounting parameter $K,$ as well as on the risk and discounting parameters together. However, no joint significance can be found (the *p*-values are 0.52, 0.72, and 0.63, respectively).

The marginal effect of the magnitude of the principal on the hyperbolic discounting parameter K is presented in Table 15.

Table 15: Treatment T0 – Marginal Effects of the Magnitude of the Principal on the Discounting Parameter in the Hyperbolic Discounting Model

Parameter	Magnitude of the Principal	Point Estimate	Standard Error	*p*-value	95% Confidence Interval		Discount rate (1 day)	Discount rate (5 years)
K	$300.00 (base category)	0.793	0.101	<0.001	0.594	0.992	121%	38%
	$100.00	0.538	0.087	<0.001	0.367	0.709	277%	50%
	$60.00	0.542	0.152	<0.001	0.243	0.841	279%	50%
	$30.00	1.186	0.239	<0.001	0.717	1.655	620%	61%
	$10.00	2.110	0.425	<0.001	1.277	2.942	1702%	73%

The estimated value of K for the highest principal of $300 is 0.79, resulting in annualized discount rates that range from 121% with a horizon of 1 day to 38% with a horizon of 5 years. The remaining rows in Table 15 refer to the estimated marginal effects on K of the other principal amounts. All estimates are significantly different from zero at all conventional levels of significance. The marginal effect on K with a principal of $100 is 0.54, resulting in a value of $K = 1.33$. The annualized discount rates implied by this new value range from 277% with a horizon of only 1 day to 50% when the horizon is extended to 5 years. Similar to the estimation of the exponential discounting, the hyperbolic model results in marginal effects on the treatment parameter K that increase as the magnitude of the principal deceases. With the lowest principal of $10, the estimated marginal effect on K has reached a sizeable increase of 211 percentage points. The new value of $K = 2.9$ implies annualized discount rates of 1702% with a horizon of 1 day, which decrease to 73% with a horizon of 5 years.

A mixture model, assuming that exponential and hyperbolic discounting (with specifications D^E and D^{H2}) are two latent DGPs, is estimated using the choice data from treatment T0. The estimation for now ignores demographic characteristics.

Table 16: Treatment T0 – Structural Estimation of Mixture Model of Exponential and Hyperbolic Discounting

Parameter	Point Estimate	Standard Error	*p*-value	95% Confidence Interval	
π^E	0.651	0.069	<0.001	0.517	0.786
π^{H2}	0.349	0.069	<0.001	0.214	0.483
r	0.662	0.014	<0.001	0.634	0.690
η	0.913	0.030	<0.001	0.854	0.971
ϕ	0.623	0.026	<0.001	0.572	0.675
δ	3.209	0.252	<0.001	2.714	3.703
K	0.382	0.081	<0.001	0.223	0.541
μ	0.140	0.006	<0.001	0.127	0.152
ν_E	0.063	0.021	0.002	0.023	0.104
ν_H2	0.518	0.141	<0.001	0.242	0.794

The p-values show that all estimates are significantly different from zero. The parameter π^E denotes the probability of exponential discounting. A Wald test clearly reveals, with $p < 0.001$, that $\pi^E$is significantly different from 0.5. Since only hyperbolic discounting is considered as an alternative DGP, the probability of this process π^{H2} equals $1 - \pi^E$. The results in Table 16 show that 65% of the choices in treatment T0 are consistent with exponential discounting and 35% are better explained by hyperbolic discounting. Separate error terms are assumed, one for each discounting process considered. The behavioral error term for exponential discounting is denoted ν_E, whereas the behavioral error term for hyperbolic discounting is given by ν_H2.

The analysis of the T0 data has illustrated some interesting methodological points. In particular, two approaches to dealing with the challenge of heterogeneous time preferences were demonstrated. A straightforward way to handle heterogeneity is to show how observable covariates affect time preferences parameters. This was implemented by conditioning choice parameters on personal and task characteristics. Another approach to heterogeneity is to account for the possibility that multiple DGPs may be present. This can be accomplished by estimating finite mixture models using MLE. Both strategies are complementary: one can have personal characteristics *and* mixture specifications.

Now that these methodological points have been established, the discussion can move on to the remaining *control treatments*. It is time to examine whether the discounting models that were introduced in section 2.1 are supported by the data. The initial focus will be on treatments T0, T1, and T3. The variation in the FED across these treatments makes it possible to examine whether the subjects' responses are consistent with the various discounting models. The MLE results are presented in Table 17.

Table 17: Treatments T0, T1, & T3 – Structural Estimation of all Discounting Models

Parameter	Point Estimate	Standard Error	*p*-value	95% Confidence Interval	
A. Exponential					
r	0.669	0.014	<0.001	0.642	0.697
η	0.912	0.030	<0.001	0.854	0.971
ϕ	0.627	0.026	<0.001	0.576	0.679
δ	1.533	0.201	<0.001	1.140	1.926
μ	0.140	0.006	<0.001	0.127	0.152
ν	0.852	0.089	<0.001	0.677	1.027
B. Hyperbolic					
r	0.669	0.014	<0.001	0.642	0.697
η	0.912	0.030	<0.001	0.854	0.971
ϕ	0.627	0.026	<0.001	0.576	0.678
K	0.998	0.090	<0.001	0.822	1.174
μ	0.140	0.006	<0.001	0.127	0.152
ν	0.846	0.089	<0.001	0.672	1.019
C. Quasi-Hyperbolic					
r	0.670	0.014	<0.001	0.642	0.697
η	0.912	0.030	<0.001	0.854	0.971
ϕ	0.627	0.026	<0.001	0.576	0.679
β	0.979	0.011	<0.001	0.958	1.001
δ	1.074	0.240	<0.001	0.604	1.545
μ	0.140	0.006	<0.001	0.127	0.152
ν	0.838	0.090	<0.001	0.662	1.014
D. Weibull					
r	0.670	0.014	<0.001	0.642	0.697
η	0.912	0.030	<0.001	0.854	0.971
ϕ	0.627	0.026	<0.001	0.576	0.679
δ	0.570	0.143	<0.001	0.289	0.851
s	1.313	0.198	<0.001	0.925	1.701
μ	0.140	0.006	<0.001	0.127	0.152
ν	0.828	0.088	<0.001	0.655	1.002

All parameters are significantly different from zero, at all conventional levels of significance.

The estimates of the parameters of the exponential and hyperbolic models are presented in Panels A and B. The estimated values of the discounting parameters δ and K are 1.53 and approximately

1.00, respectively. These estimates are lower than those obtained from treatment T0 alone. The estimated annualized discount rate for the exponential model is, of course, 153%. The discount rates implied by the hyperbolic discounting model are presented in Table 18.

Table 18: Treatments T0, T1, & T3 – Discount Rates for Several Horizons Implied by Hyperbolic Discounting

Horizon	Point Estimate	Standard Error	*p*-value	95% Confidence Interval	
1 day	1.709	0.243	<0.001	1.234	2.185
2 days	1.706	0.242	<0.001	1.232	2.179
3 days	1.702	0.241	<0.001	1.231	2.174
4 days	1.699	0.240	<0.001	1.229	2.168
5 days	1.695	0.239	<0.001	1.227	2.163
6 days	1.691	0.238	<0.001	1.225	2.157
7 days	1.688	0.237	<0.001	1.224	2.152
2 weeks	1.659	0.229	<0.001	1.210	2.108
1 month	1.608	0.216	<0.001	1.185	2.032
2 months	1.518	0.194	<0.001	1.138	1.897
3 months	1.438	0.175	<0.001	1.094	1.781
6 months	1.247	0.135	<0.001	0.983	1.511
1 year	0.998	0.090	<0.001	0.822	1.174
2 years	0.731	0.052	<0.001	0.629	0.833
5 years	0.431	0.021	<0.001	0.388	0.473

Table 18 exhibits an annualized discount rate that declines from 171% with a horizon of 1 day to 43 % when the horizon is extended to 5 years.

The estimates of the parameters from the quasi-hyperbolic model are presented in Panel C of Table 17. The quasi-hyperbolic estimate of δ is 1.07, with a 95 % CI between 0.6 and 1.55. Hence, the model proposes that the annualized discount rate, after initially imposing high fixed costs for the delayed reward, will eventually converge to 107 percent, as the horizon increases. Whether there are initial fixed costs, in form of a fixed proportion of the principal, depends on the value of the parameter β. If the estimated value of β equals 1, the model collapses to the exponential model. Indeed, the value is not significantly different from 1, as the 95 percent CI shows. This result is confirmed by a formal hypothesis test, which reveals that β falls short of being statistically different from 1 at the 5 percent level of significance (two-sided *p*-value = 0.057). This means that no significance "present bias" in the quasi-hyperbolic sense can be found in the *control treatments*. This becomes strikingly obvious when the discount rates, implied by the quasi-hyperbolic model for several horizons, ranging from 1 day to 5 years, are considered. These discount rates are constant, which makes the quasi-hyperbolic indistinguishable from exponential discounting.

The results from the Weibull specification are presented in Panel D of Table 17. The Weibull model also nests exponential discounting, in the sense that it collapses to the exponential specification when the parameter s takes a value of 1. A formal hypothesis test reveals that s is not significantly different from 1 at the 5 percent levels of significance (*p*-value = 0.11).

When the discounting models from section 2.1 are considered individually it appears that only exponential and hyperbolic discounting are well supported by the data. This demonstrates the need for some metric that can help us decide which of these two models is better.
positive.

Chapter 6 presented to popular nonparametric hypothesis tests. The Vuong and the Clarke tests can now be applied to decide which model explains the choices in treatments T0, T1, and T3 better.[17] These results can then be compared with the mixture model approach. The Vuong statistic equals -7.037. The test was set up so that a positive value would support the exponential model and a negative value supports the hyperbolic model. Hence, the Vuong test

[17] *Stata* has not integrated the two non-nested hypothesis test considered in version: Stata 13.1. The only exception is that the "vuong" option, available for zero-inflated negative binomial models. This option induces the Vuong test in this special case. Therefore, user-written tests have to be implemented, where the coding follows the test presentations in Vuong (1989) and Clarke (2003, 2007). The code used in this chapter was provided by Harrison & Rutström (2009).

favors the hyperbolic model. However, a look at the distribution of log-likelihoods in Figure 26 reveals that the log-likelihood ratios are multi-modal and sharply peaked, compared to the normal distribution.

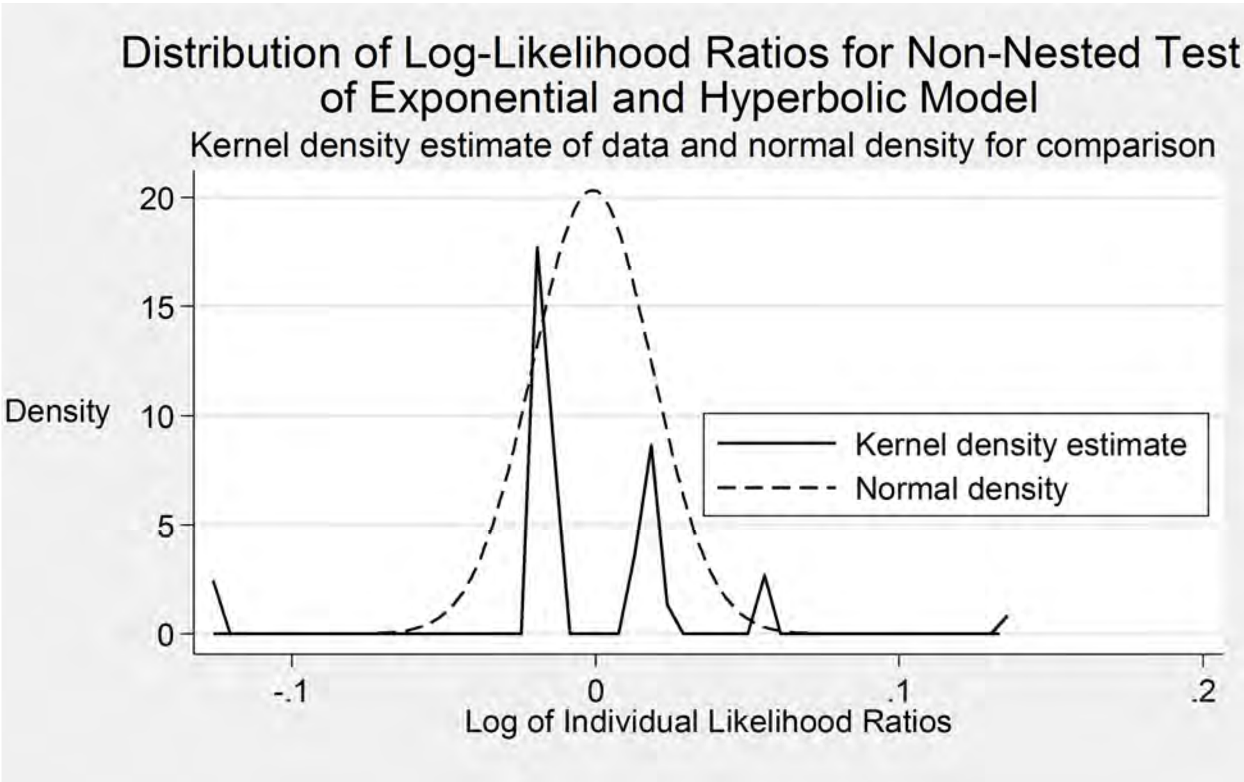

Figure 26: Distribution of Log-Likelihood Ratios, compared to a Standard Normal Distribution

Under these conditions, the Clarke test may perform better. The Clarke test is set up so that it favors the exponential model if the test statistic takes a greater value than 22539, i.e. more than half of the log-likelihood differences are positive. The actual value is 34209, so the Clarke test suggests that the exponential model explains the data better. However, this result is not statistically significant. The null-hypothesis of exponential discounting not being the better model cannot be rejected (*p*-value =1).

This situation would not be very conclusive if one and only one model had to win the comparison and be regarded as the sole DGP. However, as the discussion of mixture specifications and the application to the choices in treatment T0 demonstrated, this is not the case. Table 19 presents the estimations of four mixture models of exponential discounting and each of the non-exponential specifications.

Table 19: Treatments T0, T1, & T3 – Structural Estimation of Mixtures of Exponential and all other Discounting Models

Parameter	Point Estimate	Standard Error	*p*-value	95% Confidence Interval	
A. Exponential_Hyperbolic					
Log-Likelihood	-247967.120				
π^{E}	0.395	0.045	<0.001	0.307	0.482
π^{H2}	0.605	0.045	<0.001	0.518	0.693
r	0.662	0.014	<0.001	0.634	0.690
η	0.913	0.030	<0.001	0.854	0.971
ϕ	0.623	0.026	<0.001	0.572	0.675
δ	0.220	0.054	<0.001	0.114	0.326
K	1.545	0.083	<0.001	1.383	1.708
μ	0.140	0.006	<0.001	0.127	0.152
ν_E	0.523	0.093	<0.001	0.341	0.705
ν_H2	0.241	0.026	<0.001	0.190	0.292
B. Exponential_Quasi-Hyperbolic					
Log-Likelihood	-247961.9				
π^{E}	0.458	0.054	<0.001	0.352	0.564
π^{QH}	0.542	0.054	<0.001	0.436	0.648
r	0.662	0.014	<0.001	0.634	0.690
η	0.913	0.030	<0.001	0.854	0.971
ϕ	0.624	0.026	<0.001	0.572	0.675
δ	0.339	0.067	<0.001	0.209	0.470
δ_QH	2.452	0.659	<0.001	1.160	3.744
β	0.797	0.074	<0.001	0.652	0.942
μ	0.140	0.006	<0.001	0.127	0.152
ν_E	0.539	0.068	<0.001	0.406	0.672
ν_QH	0.285	0.108	0.008	0.074	0.496
C. Exponential_Weibull					
Log-Likelihood	-261760.740				
π^{E}	0.506	0.026	<0.001	0.456	0.557
π^{W}	0.494	0.026	<0.001	0.443	0.544
r	0.662	0.014	<0.001	0.634	0.690
η	0.913	0.030	<0.001	0.854	0.971
ϕ	0.624	0.026	<0.001	0.572	0.675
δ	0.326	0.032	<0.001	0.264	0.388
δ_wei	1.269	0.098	<0.001	1.077	1.461
s	1.191	0.100	<0.001	0.994	1.387
μ	0.140	0.006	<0.001	0.127	0.152
ν_E	0.520	0.051	<0.001	0.420	0.620
ν_Wei	0.151	0.034	<0.001	0.084	0.219

Only in the mixture model that includes exponential and hyperbolic discounting is the probability π^E of exponential discounting significantly different from 0.5 (*p*-value<0.02). With the mixtures of exponential discounting with quasi-hyperbolic and Weibull discounting, one cannot reject the hypothesis that $\pi^E = 0.5$ (*p*-values of 0.44 and 0.81, respectively). The finding that the proportion of exponential discounting choices is higher in these models than in the mixture with hyperbolic discounting deserves further attention, since it is quite possible that what appears to be the mixture of exponential and other discounting models is in fact a mixture of two exponential processes. A look at the other parameters may reveal that the alternative non-exponential specification is not significantly different from exponential discounting.

Regarding the mixture model of exponential and quasi-hyperbolic discounting, one can reject the hypothesis. The value of β is significantly less than 1. It may seem puzzling at first that the model with quasi-hyperbolic discounting as a single process collapses to an exponential model, whereas this is not the case for the quasi-hyperbolic component in the mixture model, applied to the same *control* data. However, it is perfectly possible that this occurs when two processes generate the data, but the selected model assumes the presence of only one process. For example, if the data are in fact generated by a mixture of exponential and quasi-hyperbolic discounting, but a single process quasi-hyperbolic model is estimated, the presence of choices that are consistent with exponential discounting will affect the parameter estimates. In particular, the standard errors in the single process model will be larger. This may lead to the result that β is not significantly different from 1 and, therefore, that the quasi-hyperbolic model collapses to an exponential specification. A mixture model can capture the actual process heterogeneity in the sample and prevent wrong conclusions, due to model misspecification (Harrison & Rutström 2009, 146).

The Weibull component, in the mixture model of exponential and Weibull discounting, on the other hand can be reduced to the exponential model. The parameter s is not significantly different from 1 at the 5 percent level of significance (*p*-value =0.06).

The previous tests on specific parameters revealed that only the hyperbolic and the quasi-hyperbolic component in the mixture models differed significantly from exponential discounting. Single parameter tests cannot be performed for the mixture of exponential and hyperbolic discounting, due to the non-nested nature of the two models. However, a simple test on the mixture parameter π^E reveals that, although exponential discounting is less prevalent, its probability is significantly different from zero (*p*-value = <0.001). This supports the hypothesized presence of both DGPs. Of course, the same procedure also works when mixture models are applied to cases, in which one could choose between models by running a single parameter test. The parameter π^E in the mixture of exponential and quasi-hyperbolic discounting is also significantly different from zero (*p*-value = <0.001).

To summarize, of the mixture models considered, only the mixture of exponential and hyperbolic and the mixture of exponential and quasi-hyperbolic discounting are reasonably supported by the data from treatments T0, T1, and T3. These models suggests there are indeed two decision making modes present in the data, as hypothesized by dual process models: “patient” and “deliberate” reasoning, as epitomized by exponential discounting, and “impulsive” choice behavior, as described by hyperbolic or quasi-hyperbolic discounting. Note that the significance of the quasi-hyperbolic component in the mixture with exponential discounting does not necessarily support the presence of visceral arousal as the main reason for impulsivity. Apart from alternative explanations (e.g. the quasi-hyperbolic specifications is particularly suitable to capture an initially high level of implied discount rates), visceral arousal is simply not plausible for a wide range of impulsive choice behavior. For these reasons, the test of the bundling hypothesis will focus on the marginal effect of the *bundling treatments* on the mixture probability π^E in the mixture of exponential and hyperbolic discounting. One can interpret it as support for the hypothesis, if the estimated effect is found to be significantly positive. In addition, it may be of interest to check how robust a possible effect on π^E is in the alternative mixture of exponential and quasi-hyperbolic discounting.

Before the attention can focus on the finding evidence for reward bundling, one must take one additional factor into consideration. Section 3.3 discussed that many researchers claim the existence of a *magnitude effect*, i.e. that people exhibit more patient behavior as rewards get larger. It appears that the presence of a magnitude effect in the experimental data is well supported by the descriptive results in Figures 9 and 20 and Table 5, as well as by the structural estimations presented in Tables 11 and 15. The structural analysis will now attempt a final assessment whether a genuine magnitude effect is present, before the discussion moves on to reward bundling. To accomplish this, the data from treatment T2 are included in the analysis.

Treatment T2 modifies the baseline treatment T0, by using principal amounts that are tripled. Table 20 presents the effect of the principal amounts on the discounting parameters from exponential and hyperbolic discounting.

Table 20: Treatment T2 – Marginal Effects of the Magnitude of the Principal on the Discounting Parameters in the Exponential and Hyperbolic Discounting Models

Parameter	Magnitude of the Principal	Point Estimate	Standard Error	*p*-value	95% Confidence Interval		Discount rate (1 day)	Discount rate (5 years)
A. Exponential								
δ	$900.00 (base category)	0.360	0.157	0.022	0.051	0.668	36%	36%
	$300.00	0.910	0.275	0.001	0.371	1.449	127%	127%
	$180.00	1.275	0.425	0.003	0.442	2.108	163%	163%
	$90.00	2.470	0.777	0.001	0.946	3.993	283%	283%
	$30.00	7.131	3.135	0.023	0.986	13.276	749%	749%
B. Hyperbolic								
K	$900.00 (base category)	0.311	0.119	0.009	0.078	0.543	36%	21%
	$300.00	0.526	0.064	<0.001	0.401	0.651	131%	19%
	$180.00	0.684	0.075	<0.001	0.536	0.831	170%	43%
	$90.00	1.089	0.105	<0.001	0.882	1.295	304%	52%
	$30.00	1.994	0.190	<0.001	1.622	2.366	895%	66%

In the case of exponential discounting the estimated annualized discount rate δ with the highest principal amount of $900 is 36%. This value remains constant, of course, regardless of the horizon. The remaining rows of panel A in Table 20 estimate the value of δ relative to the base category of $900. All estimates are significantly different from zero at the 5 percent level of significance. The estimated marginal effect increases, as the magnitude of the principal decreases. A principal of $300 results in marginal effect that increases the value of δ by 91 percentage points, resulting in an annualized discount rate of 127% for all horizons. The marginal effect, relative to the base category, increases to 7.13, when the principal is lowered to $30, and the resulting constant discount rate is 749%.

If hyperbolic discounting is assumed, panel B of Table 20 shows that K takes an estimated value of 0.31 for the base category of $900. The discount rates implied by this estimate range from 36% with a horizon of 1 day to 21% with a horizon of 5 years. The other rows exhibit the marginal effects, relative to the base category, when the principal is lowered. All estimates are significantly different from zero at all conventional levels of significance. When the principal is lowered from $900 to $300, the estimate of K increases by 0.53: This resulting value of $K = 0.84$ implies an annualized discount rate that ranges from 131% with a horizon of 1 day to 19% with a horizon of 5 years. The lowest principal in treatment T2 of $30 results in K increasing by 1.99, relative to the base category. The new value of $K = 2.31$ implies discount rates from 895% with a horizon of 1 day to 66% with a horizon of 5 years.

Similarly to treatment T0, the data from treatment T2 clearly suggest an effect of the magnitude of the principal on the discounting parameters. It is instructive to compare the estimated parameters from both treatments in greater detail. If the higher principals in treatment T2 result in inferred discounting parameters that are significantly lower than those inferred from T0 choices, this would strongly support the presence of a magnitude effect. Table 21 presents this comparison.

Table 21: Treatments T0 and T2 – Exponential and Hyperbolic Discounting Parameters

Parameter	Point Estimate	Standard Error	*p*-value	95% Confidence Interval	
A. Treatment T0					
δ	2.120	0.335	<0.001	1.464	2.777
K	1.189	0.117	<0.001	0.960	1.419
B. Treatment T2					
δ	1.129	0.221	<0.001	0.696	1.562
K	0.778	0.110	<0.001	0.563	0.994

All estimates are significantly different from zero. The results in panel A have already been presented in Tables 9 and 12. They are included again to facilitate the comparison. A first look at the equivalent results in panel B of Table 21 suggest that the parameters inferred from the choices in treatment T2 are much lower. Indeed, a formal test confirms that the null-hypothesis of equal δ estimates in treatments T0 and T2 can clearly be rejected (*p*-value = 0.009). Similarly, a formal test of the null hypothesis that the hyperbolic discounting parameter K is equal in both treatments can also be rejected with high confidence (both *p*-value = 0.007).

As detailed in subsection 2.2.4, behavior that is consistent with a magnitude effect may actually not be a response to the magnitude of rewards. Instead, it may be the result of rounding behavior. Given the theoretical reasons for this argument, rounding cannot plausibly explain the difference in discounting behavior between the choices in treatments T0 and T2. The purpose of rounding is to lower the cognitive burden of comparing similar amounts. However, only small principals result in similar LL amounts, when the annualized interest rates are applied. The large principals in treatment T2 lead to dissimilar LL amounts. Comparing SS and LL amounts should, therefore, not pose particular difficulties for subjects, and no rounding is needed to simplify the task. One can conclude that the increased patience that occurs with higher principals is the result of a genuine magnitude effect, as already suggested by the descriptive probit model that investigated the marginal effect of SCALE (Figure 19 and Table 5). The presence of a magnitude effect has important consequences for the subsequent analysis of the bundling data. One must be able to distinguish possible reactions to the magnitude of the principal from responses to multiple decision pairs.

The analysis can now proceed and take all data into consideration. A mixture model, assuming the co-existence of exponential and hyperbolic discounting, is applied to the data from all *treatments*. The results are presented in Table 22.

Table 22: Bundling treatments – Structural Estimation of Mixture Model of Exponential and Hyperbolic Discounting

Parameter	Point Estimate	Standard Error	*p*-value	95% Confidence Interval	
π^{E}	0.445	0.033	<0.001	0.381	0.509
π^{H2}	0.555	0.033	<0.001	0.491	0.619
r	0.665	0.014	<0.001	0.637	0.692
η	0.913	0.030	<0.001	0.854	0.971
ϕ	0.625	0.026	<0.001	0.574	0.676
δ	0.246	0.037	<0.001	0.173	0.319
K	1.624	0.108	<0.001	1.413	1.835
μ	0.140	0.006	<0.001	0.127	0.152
ν_*E*	0.535	0.064	<0.001	0.409	0.662
ν_*H2*	0.223	0.016	<0.001	0.192	0.253

All estimates are statistically highly significant, in the sense of being different from zero. The respective estimates of the discounting parameters of the exponential and hyperbolic model, δ and K, are 0.25 and 1.62, respectively. These values are comparatively close to the equivalent estimates found with only the control data, in panel A of Table 19. The discount rates, implied by this estimated value of K are presented in Table 23.

Table 23: Bundling treatments – Discount Rates for Several Horizons Implied by the Hyperbolic Discounting Parameter *K*

Horizon	Point Estimate	Standard Error	*p*-value	95% Confidence Interval	
1 day	4.056	0.542	<0.001	2.994	5.119
2 days	4.038	0.538	<0.001	2.984	5.092
3 days	4.020	0.534	<0.001	2.974	5.066
4 days	4.003	0.529	<0.001	2.965	5.040
5 days	3.985	0.525	<0.001	2.956	5.014
6 days	3.968	0.521	<0.001	2.946	4.989
7 days	3.950	0.517	<0.001	2.937	4.964
2 weeks	3.814	0.486	<0.001	2.863	4.766
1 month	3.587	0.435	<0.001	2.734	4.440
2 months	3.210	0.357	<0.001	2.511	3.909
3 months	2.909	0.299	<0.001	2.322	3.495
6 months	2.284	0.195	<0.001	1.901	2.666
1 year	1.624	0.108	<0.001	1.413	1.835
2 years	1.061	0.052	<0.001	0.959	1.164
5 years	0.556	0.018	<0.001	0.520	0.592

The annualized discount rates, implied by the estimate of K range from 406% for a horizon of 1 day to 56% for horizon of 5 years.

Of particular interest is the value of π^E of 0.45 in Table 22. This means that about 45% of all choices are consistent with exponential discounting, whereas 55% of the choices are consistent with hyperbolic discounting. This probability of exponential discounting is higher in all treatments, compared to the 39.5% in the control treatments, which are displayed of Panel A of Table 19. This difference suggests some tentative support for the bundling hypothesis. However, in order to decide whether the presentation with multiple decisions is truly responsible for the observed increase in the probability of exponential discounting, one must examine which factors actually affect the mixture parameters.

The analysis of all data must be extended by including estimating the mixture model of exponential and hyperbolic, while controlling for the full set of demographic and treatment covariates discussed so far. Of particular interest is the effect of BUNDLED on the probability of exponential discounting. A comparison of the probability of exponential discounting in the *control treatments* and in the *bundling treatments* is presented in Figure 27.

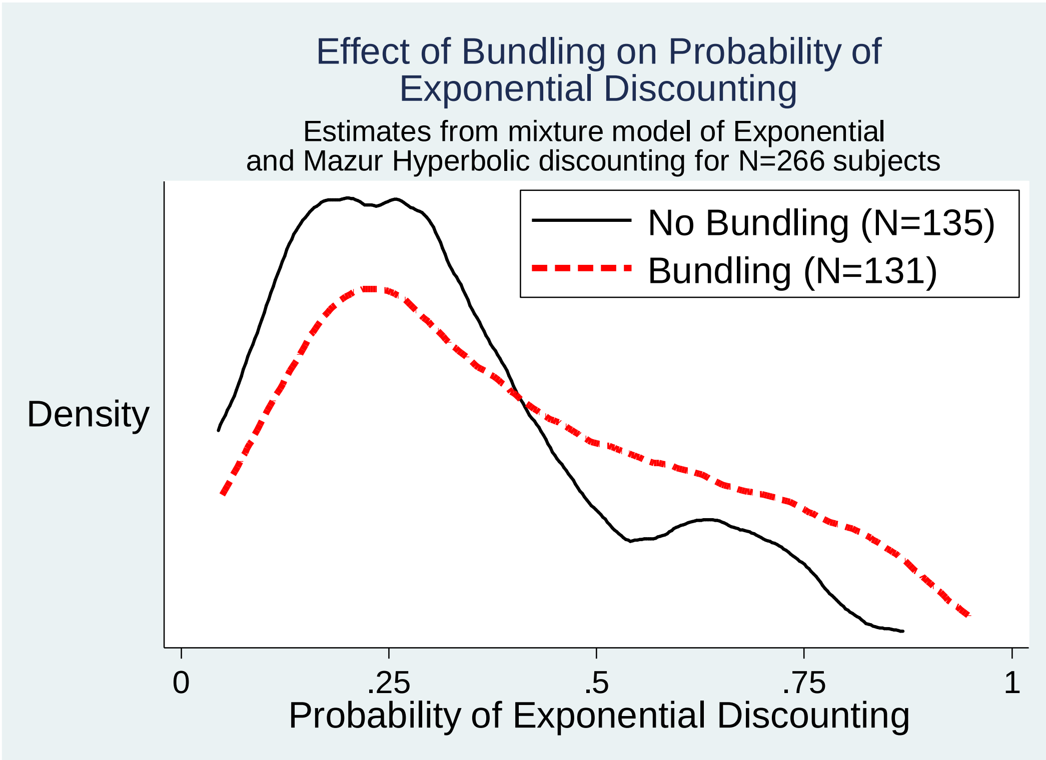

Figure 27: Total Effect of Bundling on the Probability of Exponential Discounting in mixture of Exponential and Hyperbolic Discounting

Figure 27 shows that a low probability of exponential discounting is clearly more likely in the control treatments than in the bundling treatments. A higher probability of exponential discounting becomes more likely in the bundling treatments. This is additional tentative support for the bundling hypothesis, but it is still not sufficient. The reason is that Figure 27 illustrates the *total effect* of bundling. In order to evaluate the pure effect of BUNDLED, one must estimate the *marginal effect*, holding all treatment and demographic covariates.

Table 24 presents the marginal effects of several characteristics of decision makers and choice situation on the core parameters of an exponential and hyperbolic mixture model.

Table 24: Bundling treatments – Marginal Effects of Demographic and Treatment Covariates on Core Parameters in Mixture Model of Exponential and Hyperbolic Discounting

Parameter	Variable	Point Estimate	Standard Error	*p*-value	95% Confidence Interval	
r	FEMALE	0.017	0.031	0.590	-0.044	0.077
	BLACK	0.003	0.031	0.936	-0.059	0.064
	ASIAN	-0.006	0.049	0.900	-0.103	0.091
	GPAHI	-0.005	0.030	0.873	-0.064	0.054
	BUSINESS	-0.024	0.043	0.576	-0.109	0.061
η	FEMALE	0.099	0.071	0.164	-0.040	0.238
	BLACK	-0.098	0.086	0.255	-0.266	0.070
	ASIAN	-0.127	0.097	0.190	-0.316	0.063
	GPAHI	-0.053	0.065	0.411	-0.180	0.074
	BUSINESS	-0.067	0.077	0.387	-0.218	0.084
ϕ	FEMALE	-0.044	0.049	0.360	-0.140	0.051
	BLACK	0.097	0.067	0.152	-0.035	0.229
	ASIAN	0.161	0.096	0.092	-0.027	0.349
	GPAHI	-0.012	0.053	0.818	-0.116	0.091
	BUSINESS	-0.067	0.066	0.315	-0.197	0.063
δ	BUNDLED	0.165	0.088	0.060	-0.007	0.337
	FEMALE	0.139	0.090	0.122	-0.037	0.315
	BLACK	-0.022	0.061	0.717	-0.142	0.097
	ASIAN	0.008	0.075	0.914	-0.140	0.156
	GPAHI	-0.064	0.061	0.296	-0.184	0.056
	BUSINESS	0.009	0.040	0.825	-0.069	0.087
K	BUNDLED	0.322	0.156	0.038	0.017	0.627
	FEMALE	-0.030	0.197	0.879	-0.416	0.356
	BLACK	0.017	0.214	0.938	-0.403	0.436
	ASIAN	-0.095	0.298	0.749	-0.679	0.489
	GPAHI	-0.120	0.179	0.502	-0.471	0.231
	BUSINESS	0.014	0.213	0.947	-0.404	0.432
π^E	BUNDLED	0.307	0.130	0.019	0.051	0.562
	FEMALE	-0.014	0.090	0.875	-0.191	0.163
	BLACK	-0.075	0.092	0.419	-0.255	0.106
	ASIAN	0.211	0.161	0.189	-0.104	0.526
	GPAHI	-0.045	0.087	0.606	*-0.214*	0.125
	BUSINESS	0.121	0.087	0.163	-0.049	0.291
	FED	0.006	0.003	0.071	-0.001	0.013
	SCALE	0.080	0.054	0.135	-0.025	0.185
	FORCED	-0.330	0.084	<0.001	-0.494	-0.166

SECOND	-0.235	0.092	0.011	-0.416	-0.054
THIRD	-0.368	0.137	0.007	-0.637	-0.099
$10.00	-0.453	0.163	0.005	-0.774	-0.133
$30.00	-0.392	0.125	0.002	-0.636	-0.147
$60.00	-0.279	0.076	<0.001	-0.428	-0.129
$100.00	-0.302	0.082	<0.001	-0.464	-0.141

Most importantly, the results indicate a positive marginal effect of BUNDLED on the probability of exponential discounting. Table 24 shows a 31 percentage point increase in the choices that are consistent with exponential discounting, compared to the choices that are better characterized by hyperbolic discounting, as the effect of introducing bundling treatments. The effect is statistically significant in the sense of being different from zero (p-value = 0.02). This is strong support for the reward bundling hypothesis.

Lower positive effects of BUNDLED on discounting parameters δ and K are also found. However, only the effect of BUNDLED on K is statistically different from zero at the 5 percent level. A test of the joint hypothesis that the effect of BUNDLED on the parameters, K, and π^E is zero can be rejected with high confidence (p-value =0.005).[18] The lack of individually significant effects of the demographic characteristics on the core parameters is confirmed when all data are considered.

Somewhat surprising is the finding of a negative effect of FORCED on the probability of exponential discounting. Table 24 shows that FORCED decreases the probability of exponential discounting by 33 percentage points, an effect that is significantly different from zero. The constraint, imposed by the first decision on the remaining decisions in treatments T5 and T7,

[18] These key results essentially hold when a mixture of exponential and quasi-hyperbolic discounting is used instead. In this case, BUNDLED increases the probability of choices consistent with exponential discounting by 30 percentage points. This strong marginal effect is significantly different from zero (two-sided p-value = 0.019). The joint significance of BUNDLED on parameters, δ, δ_{QH}, β, and π^E is also significantly different from zero at the 5 percent level of significance (p-value = 0.048).

could function as a pre-commitment device to support impulse control. In this case, one would expect that FORCED increases the probability of exponential discounting. However, the actual result lends additional support to the notion that the increase in the fraction of exponential choices, resulting from the inclusion of bundling treatments, is indeed the consequence of reward bundling. Reward bundling in human subjects is a psychological process prompted by the presentation of each decision as part of a series of comparable decisions. If reward bundling is truly responsible for the increase in choices that are consistent with exponential discounting, then one would expect this increase to occur in all tasks with multiple decisions. This is indeed exactly what the data reveal. A possible separate effect from pre-commitment due to an experimentally induced constraint of the FORCED condition must not be confused with the genuine bundling effect.

No statistically significant effect of SCALE on the mixture parameter π^E is found. This is additional support for reward bundling, since it rules out that the observed increase in exponential discounting is merely the result of the bundling treatments multiplying the payouts. The negative effects of SECOND and THIRD, on the other hand, are significantly different from zero. This is consistent with the finding from the descriptive probit models that the positive effect of including the bundling treatments is limited to the first choices.

Strong magnitude effects from the value of the principal on the mixture parameter π^E are found. They indicate that the probability of exponential discounting decreases when the highest principal is lowered.

9. Conclusion

Reward bundling has long been recognized as a possible process behind impulse control. Choice patterns consistent with this hypothesis have repeatedly been found with animal subjects (Anslie 1974, 1975; Ainslie & Monterosso 2003). However, no attempt has been made to test for the bundling phenomenon with human subjects, in a way that satisfies the methodological standards of experimental economics.

The experiment reported in this chapter closes this gap. It presented undergraduate subjects with a series of experimental treatments. Each treatment condition varied important aspects of the choice task which might affect the patience that subjects display when choosing between time dated rewards. This procedure allowed the identification of subjects discounting behavior and controlling for several factors, other than the number of decisions, which affect their observable choices.

The resulting data are analyzed using descriptive and structural econometric procedures which are well-established in the binary choice literature. Descriptive models include local polynomial regression and probit models which relate the tendency to display more patient choice behavior to several task characteristics. The key aspect of the decision situation is the presence of the *bundling condition*, where each individual decision is presented as part of a series of comparable decisions with different payout dates. The emphasis of the econometric investigation is placed on a comprehensive structural analysis.

Full-information maximum likelihood estimation is employed to estimate several common functional forms that have emerged from the literature on discounting behavior (Andersen et.al 2014). Recent findings that inferred discount rates are highly sensitive to the curvature of the utility functions and must, therefore, be jointly estimated with risk attitudes are taken into consideration (Anderson et.al. 2008). The analysis does, once again, confirm this sensitivity and the importance of joint estimation for an accurate assessment of discount rates.

A restrictive declaration that one and only one discounting model outperforms all competing accounts is usually an oversimplification, which may not do justice to the complexity of human decision making. Finite mixture models can provide valuable insights into what aspects of the choice situation affect the relative support of one theory over another. In the present context they allow evaluating when more patient and deliberate choice behavior, exemplified by exponential discounting, becomes more likely than impulsive behavior, described by a hyperbolic discounting specification. This analytical framework offers a natural way to test whether the subjects' behavior is consistent with reward bundling.

The results offer strong support for the reward bundling hypothesis. Descriptive and structural methods alike find that the tendency to exhibit more patient choice behavior increases drastically when each decision is presented as part of a series of time dated decisions. Other factors that may explain this observation, such as scale effects or imposing "forced" links of individual decisions, are carefully controlled for.

The presence of a significant magnitude effect is found in the data. This result is noteworthy, given that the findings of magnitude effects that are frequently asserted in the literature, can often be explained by rounding operations that subjects may employ to ease the cognitive burden of the choice task. The magnitude effect found in the present data is robust to this objection, since the high principals employed do not result in later amounts that are very similar to the principal amounts.

References

Ainslie, G. W. (1974). Impulse Control in Pigeons. *Journal of the experimental analysis of behavior*, 21(3), 485-489.

Ainslie, G. (1975). Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychological Bulletin*, 82(4), 463.

Ainslie, G. (1992). *Picoeconomics: The strategic interaction of successive motivational states within the person*. Cambridge: Cambridge University Press.

Ainslie, G. (2001). *Breakdown of will*. Cambridge University Press.

Ainslie, G. (2012). Pure hyperbolic discount curves predict "eyes open" self-control. *Theory and Decision*, 73(1), 3-34.

Ainslie, G.W. and Monterosso, J.R. (2003) Building blocks of self-control: Increased tolerance for delay with bundled rewards. *Journal of the Experimental Analysis of Behavior*,79, 83-94.

Akaike, H. (1973). Information Theory and an Extension of the Maximum Likelihood Principle. In *Second International Symposium on Information Theory*. Akademinai Kiado, 267-281.

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

Andersen, S., Harrison, G. W., Lau, M. I., and Rutström, E. E. (2010). Behavioral econometrics for psychologists. *Journal of Economic Psychology*, 31, 553-576.

Andersen, S., Harrison, G.W., Lau, M. I., and Rutström, E. E. (2011). *Multiattribute Utility Theory, Intertemporal Utility and Correlation Aversion*. *Working Paper 2011-04*, Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University.

Andersen, S., Harrison, G.W., Lau, M. I., and Rutström, E. E. (2013). Discounting Behaviour and the Magnitude Effect: Evidence from a Field Experiment in Denmark. *Economica*, 80(320), 670-697.

Andersen, S., Harrison, G. W., Lau, M. I., & Rutström, E. E. (2014a). Dual criteria decisions. *Journal of Economic Psychology*, 41, 101-113.

Andersen, S., Harrison, G.W., Lau, M. I., and Rutström, E. E. (2014b). Discounting Behavior: A Reconsideration. *European Economic Review*, 71, 15-33.

Andersen, S., Cox, J.C., Harrison, G. W., Lau, M., Rutström, E. E., and Sadiraj, V. (2011). Asset Integration and Attitudes to Risk: Theory and Evidence. *Working Paper 2011-17*. Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University.

Andreoni, J., & Sprenger, C. (2012). Risk Preferences Are Not Time Preferences. *The American Economic Review,* 102(7), 3357-3376.

Barrett, L. F., Tugade, M. M., & Engle, R. W. (2004). Individual Differences in Working Memory Capacity and Dual-Process Theories of the Mind. *Psychological Bulletin*, 130(4), 553-573.

Benhabib, J., Bisin, A., & Schotter, A. (2010). Present-bias, quasi-hyperbolic discounting, and fixed costs. *Games and Economic Behavior*, 69(2), 205-223.

Böhm-Bawerk, E. v. (1889), 1970. *Capital and Interest.* South Holland: Libertarian Press.

Chung, S. H., & Herrnstein, R. J. (1967). Choice and Delay of Reinforcement. *Journal of the Experimental Analysis of Behavior*, 10(1), 67-74.

Clarke, K. A. (2003). Nonparametric model discrimination in international relations. *Journal of Conflict Resolution*, 47(1), 72-93.

Clarke, K. A. (2007). A Simple Distribution-Free Test for Nonnested Model Selection. *Political Analysis*, 15(3), 347-363.

Coller, M., Harrison, G. W., & Rutström, E. E. (2012). Latent Process Heterogeneity in Discounting Behavior. *Oxford Economic Papers*, 64(2), 375-391.

Coller, M., & Williams, M. B. (1999). Eliciting individual discount rates. *Experimental Economics*, 2(2), 107-127.

Cox, J. C., & Sadiraj, V. (2006). Small-and large-stakes risk aversion: Implications of concavity calibration for decision theory. *Games and Economic Behavior*, 56(1), 45-60.

Edwards, W. (1962). Subjective Probabilities inferred from Decisions. *Psychological Review,* 69(2), 109-135.

Elster, J. (1979). *Ulysses and the sirens: Studies in rationality and irrationality*. Cambridge: Cambridge University Press.

Evans, J. S. B. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology*, 59, 255-278.

Fan, J., & Gijbels, I. (1996). *Local Polynomial Modelling and Its Applications*. London: Chapman & Hall.

Fechner, G.T. (1860). Elemente der Psychophysik. Leipzig, Germany: Breitkopf and Härtel.

Fisher, I. (1930). *The Theory of Interest.* NY: Macmillan.

Frederick, S., Loewenstein, G., & O'Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40(2), 351-401.

Fudenberg, D., & Levine, D. K. (2012). Timing and Self-Control. *Econometrica*, 80(1), 1-42.

Fudenberg, D., & Levine, D. K. (2006). A Dual-Self Model of Impulse Control. *American Economic Review*, 96, 1449-1476.

Gasser, T., & Müller, H.-G. (1979). Kernel estimation of regression functions. In *Smoothing Techniques for Curve Estimation, Lecture Notes in Mathematics*, ed. Gasser, T. and Rosenblatt, M., 23–68. New York: Springer.

Georgescu-Roegen, N. (1958). Threshold in Choice and the Theory of Demand. *Econometrica*, 26(1), 157-168.

Glimcher, P., Kable, J., and Louie, K. (2007). Neuroeconomic studies of impulsivity: Now or just as soon as possible? *American Economic Review*, 97, 1-6.

Harrison, G. W., & Rutström, E. E. (2008). Risk Aversion in the Laboratory, in Cox, J.C. & Harrison, G.W. (Eds.), *Risk Aversion in Experiments* (Vol. 12). Emerald Group Publishing, Bingley, UK.

Harrison, G. W., & Rutström, E. E. (2009). Expected Utility Theory and Prospect Theory: One Wedding and a Decent Funeral. *Experimental Economics*, 12(2), 133-158.

Harrison, G.W. and List, J.A. (2004). Field Experiments. *Journal of Economic Literature*, 42(4), 1013–1059.

Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior*, 4, 267-272.

Hey, J. D., & Orme, C. (1994). Investigating generalizations of expected utility theory using experimental data. *Econometrica*, 62(6), 1291–1326.

Jevons, H. S. (1905). *Essays on Economics.* London: Macmillan.

Jevons, W. S. (1888). *The Theory of Political Economy.* London: Macmillan.

Johnson, J. G., & Busemeyer, J. R. (2010). Decision making under risk and uncertainty. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1(5), 736-749.

Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2), 263-291.

Kirby, K. N., & Guastello, B. (2001). Making choices in anticipation of similar future choices can increase self-control. *Journal of Experimental Psychology: Applied*, 7(2), 154-164.

Kullback, S. (1959). *Information Theory and Statistics.* New York: Wiley.

Laibson, D. (1997). Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics,* 62, 443-479.

Laury, S. K., McInnes, M. M., & Swarthout, J. T. (2012). Avoiding the curves: Direct elicitation of time preferences. *Journal of Risk and Uncertainty*, 44(3), 181-217.

Loewenstein, G. (1996). Out of control: Visceral influences on behavior. *Organizational Behavior and Human Decision Processes,* 35, 272–292.

Loewenstein, G., & Prelec, D. (1992). Anomalies in Intertemporal Choice: Evidence and Interpretation. *Quarterly Journal of Economics*, 107(2), 573-597.

Loomes, G. (2005). Modelling the Stochastic Component of Behaviour in Experiments: Some Issues for the Interpretation of Data. *Experimental Economics*, 8(4), 301-323.

Lopes, L. L. (1987). Between Hope and Fear: The Psychology of Risk. *Advances in Experimental Social Psychology*, 20(3), 255-295

Lopes, L. L., & Oden, G. C. (1999). The Role of Aspiration Level in Risky Choice: A Comparison of Cumulative Prospect Theory and SP/A Theory. *Journal of Mathematical Psychology*, 43(2), 286-313.

Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. In M. L. Commons, J. E. Mazur, J. A. Nevin, & H. Rachlin (Eds.), *Quantitative analyses of behavior V: The effect of delay and of intervening events on reinforcement value*. Hillsdale: Erlbaum.

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

McClure, S. M., Laibson, D. I., Loewenstein, G., & Cohen, J. D. (2004). The grasshopper and the ant: Separate neural systems value immediate and delayed monetary rewards. *Science, 306*, 503–507.

Phelps, E. S. & Pollack, R.A. (1968). On second-best national saving and game-equilibrium growth. *Review of Economic Studies, 35*, 185–199.

Prelec, D. (1998). The Probability Weighting Function. *Econometrica*, 66(3), 497-528.

Quiggin, J. (1982). A Theory of Anticipated Utility. *Journal of Economic Behavior & Organization*, 3(4), 323–343.

Rachlin, H. & Green, L. (1972). Commitment, choice, and self-control. *Journal of the Experimental Analysis of Behavior*, 17,15-22.

Rae, J. (1834). *The Sociological Theory of Capital (reprint 1834 ed.).* London: Macmillan.

Read, D. (2001). Is time-discounting hyperbolic or subadditive? *Journal of Risk and Uncertainty*, 23(1), 5-32.

Rubinstein, A. (1988). Similarity and decision-making under risk (Is there a utility theory resolution to the Allais paradox?). *Journal of Economic Theory*, 46(1), 145-153.

Saha, A. (1993). Expo-power utility: A 'flexible' form for absolute and relative risk aversion. *American Journal of Agricultural Economics*, 75(4), 905-913.

Samuelson, P. A. (1937). A note on measurement of utility. *The Review of Economic Studies*, 4(2), 155-161.

Scholten, M., & Read, D. (2010). The Psychology of Intertemporal Tradeoffs. *Psychological Review*, 117(3), 925-944.

Schwarz, G. (1978). Estimating the Dimension of a Model. *Annals of Statistics*,6,461-64.

Starmer, C. (2000). Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk. *Journal of Economic Literature*, 332-382.

Strotz, R. H. (1956). Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies*, 23, 166-180.

Thaler, R. (1981). Some empirical evidence on dynamic inconsistency. *Economics Letters*, 8(3), 201-207.

Train, K. (2009). *Discrete Choice Methods with Simulation*. 2nd edition. Cambridge: Cambridge University Press.

Tversky, A. (1977). Features of Similarity. *Psychological Review*, 84(4), 327-352.

Tversky, A. & Kahneman, D. (1992). Advances in Prospect Theory: Cumulative Representations of Uncertainty. *Journal of Risk & Uncertainty*, 5, 297-323.

Vickrey, W. (1945). Measuring Marginal Utility by Reactions to Risk. *Econometrica*, 13(4), 319-333.

Vuong, Quang. 1989. Likelihood Ratio Tests for Model Selection and Nonnested Hypotheses. *Econometrica*, 57, 307-33.

Watson, G. S. (1964). Smooth regression analysis. *Sankhyā Series* A 26: 359–372.

Weber, E.H. (1834). *De pulsu, resor tione, auditu et tactu annotatin es anatomicae et physiologicae*. Leipzig, Germany: Köhler

Wilcox, N. T. (2008). Stochastic Models for Binary Discrete Choice under Risk: A Critical Primer and Econometric Comparison. In: J. Cox & G. W. Harrison (Eds), *Risk Aversion in Experiments* (Vol. 12). Bingley, UK: Emerald, Research in Experimental Economics.

Wilcox, N. T. (2011). 'Stochastically more risk averse': A contextual theory of stochastic concrete choice under risk. *Journal of Econometrics*,162, 89–104.

Appendix: Instructions

A. Instructions: Treatment T0

B_T0

Sooner versus Later Payments

In this task you will make a number of choices between receiving an amount of money on a "sooner" date or a different amount of money on a "later" date. The sooner date will always be today, while the later date will be some weeks from today. An example of a decision screen is shown below. The dates and amounts in your task will differ. You will make all decisions on a computer.

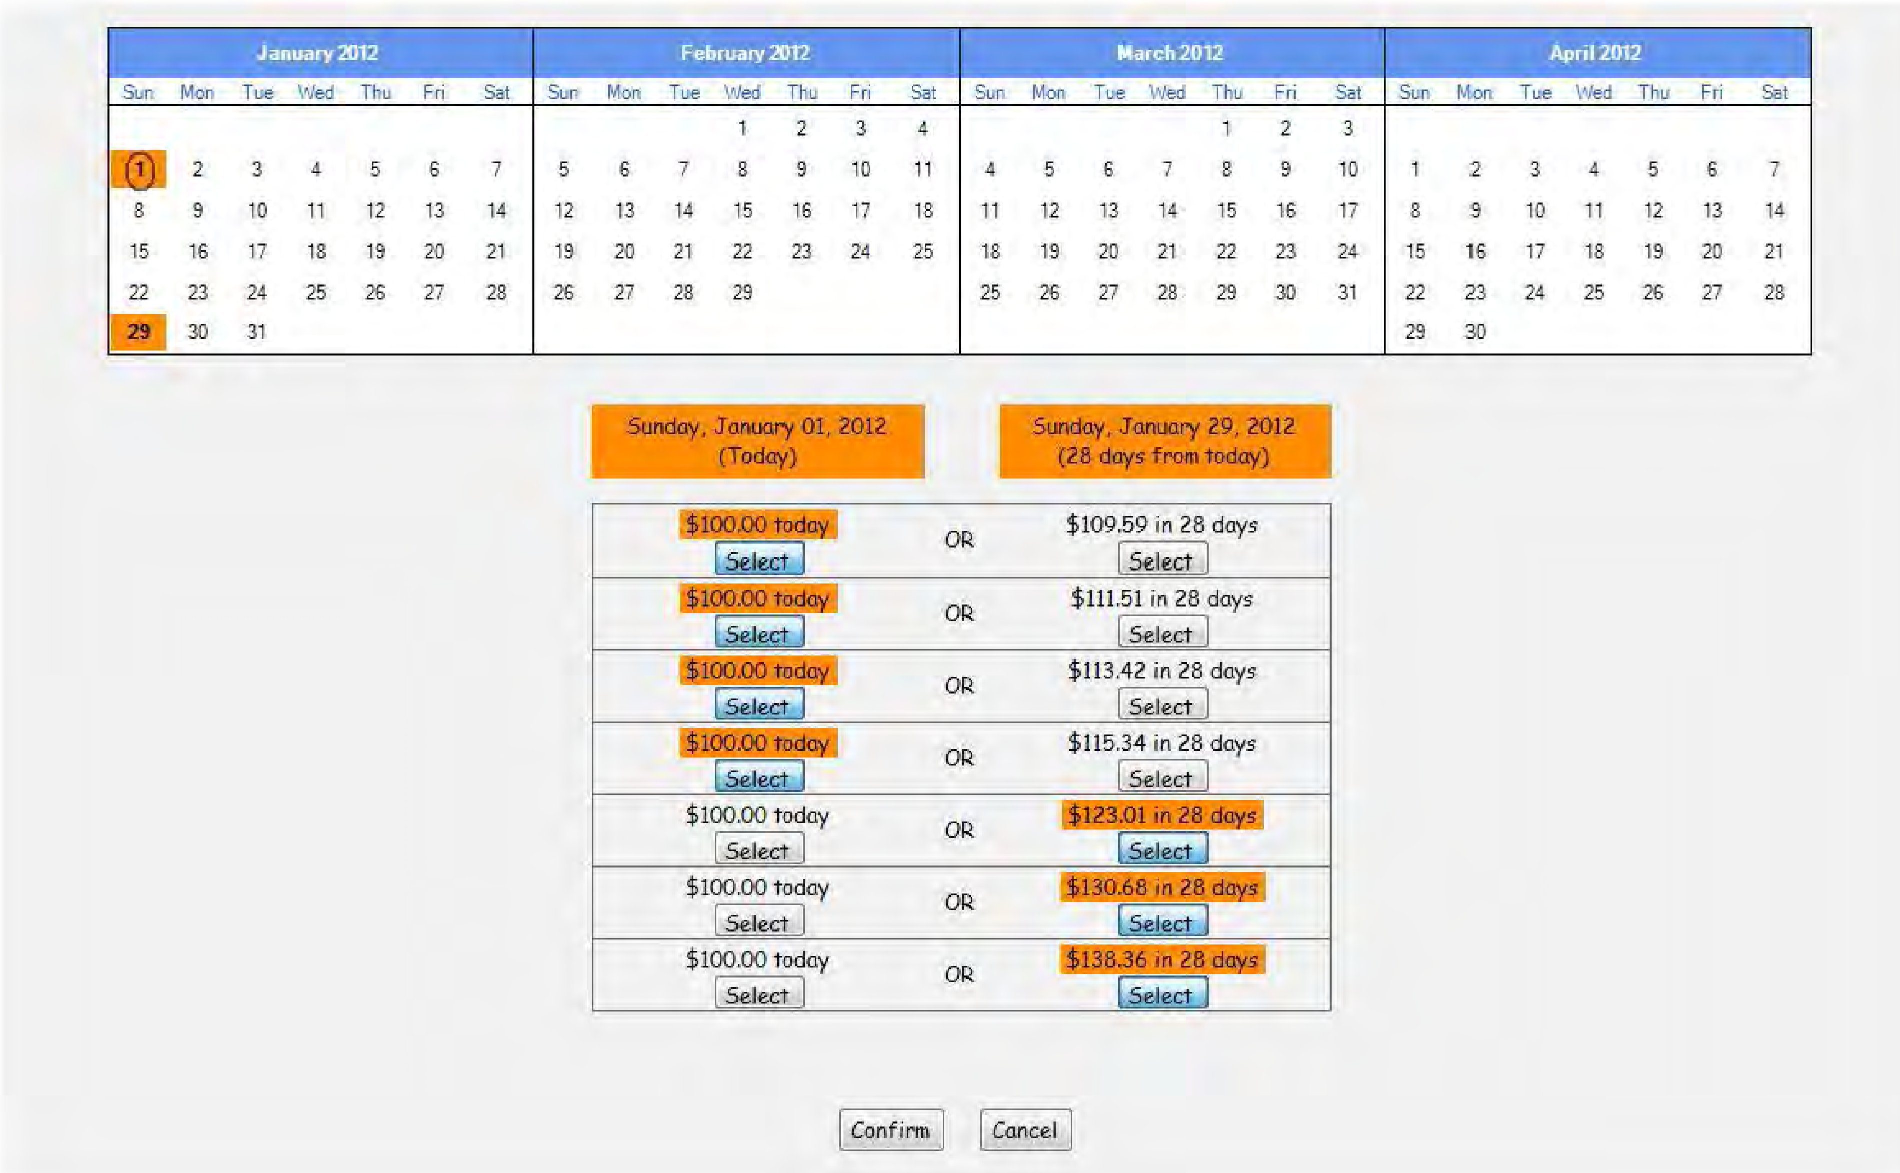
January 2012
February 2012
March 2012
April 2012
Sun Mon Tue Wed Thu Fri Sat
Sunday, January 01, 2012 (Today)
Sunday, January 29, 2012 (28 days from today)
$100.00 today OR $109.59 in 28 days
$100.00 today OR $111.51 in 28 days
$100.00 today OR $113.42 in 28 days
$100.00 today OR $115.34 in 28 days
$100.00 today OR $123.01 in 28 days
$100.00 today OR $130.68 in 28 days
$100.00 today OR $138.36 in 28 days
Select
Confirm
Cancel

This screen shows seven decisions. Each decision is presented on a different row. All decisions have the same format. For the purpose of explaining this task, assume for the moment that today is January 1, 2012. In the calendar, today's date is enclosed in a black circle. Let's look at the first decision in the example above (the one on the first decision row). The sooner choice pays $100 today, January 1, 2012 in the example, and the later choice pays $109.59 in twenty-eight days from today. You choose between these two options by clicking the button under the option you prefer.

We will present you with ten of these decision screens, with each screen having seven choices for you to make. You must make all seven choices on each decision screen before moving to the next decision screen. While on a single decision screen, the only difference between decisions is that the dollar amounts of the future payment will change. However, different decision screens will have different dollar amounts and future payment dates. So, you should make sure to pay attention to both the changing dollar amounts and changing dates as you make your decisions.

After you have worked through all of the decisions, we will select one of your ten decision screens by rolling a 10-sided die. Then we will roll a 10-sided die again, until a number between 1 and 7 comes up, to pick one decision row on that screen. When you make your choices you will not know which decision will be selected for payment. You should therefore treat each decision as if it might actually count for payment.

Once the decision screen and row are selected, we will look at the specific choice you made and the payout for that choice. If the payout is smaller (less than $100), you will actually be paid this amount. However, if the payout is larger ($100 or greater), you will roll a 10-sided die to determine whether or not you are actually paid this amount. If you roll a 1, you will actually be paid the amount, and on the date that you chose to receive it. If you roll a number other than 1, you will earn nothing in this task. This roll will be at the end of the experiment, and in private, when you are being paid.

For instance, suppose the decision screen in the above example was selected and you preferred the sooner date in the first four rows and the later date in the last three rows, as shown above. You would then roll a 10-sided die, until a number between 1 and 7 comes up, to select the row. Suppose the outcome was 5 and the fifth row is selected, where you preferred a payment of $123.01 in twenty-eight days. Since this payment is $100 or greater, you would roll the 10-sided die again, at the end of the experiment. If the outcome of your die-roll is 1, you will be paid $123.01 in twenty-eight days. However, if the outcome is 3, or anything other than 1, you will get nothing.

You will receive the money on the date stated in your preferred option. We will pay you using *PayPal*, which is an online payment service. We will explain more about *PayPal* in a few minutes. If you receive some money to be paid in the future you will also receive a written confirmation from Professor Harrison which guarantees that the money is to be paid to you on that date.

The money you receive from these choices is in addition to the show-up fee of $5, which is paid out at the end of the experiment as cash.

B. Instructions: Treatment T1

B_T1

Sooner versus Later Payments

In this task you will make a number of choices between receiving an amount of money on a "sooner" date or a different amount of money on a "later" date. The sooner date will always be tomorrow, while the later date will be some weeks from today. An example of a decision screen is shown below. The dates and amounts in your task will differ. You will make all decisions on a computer.

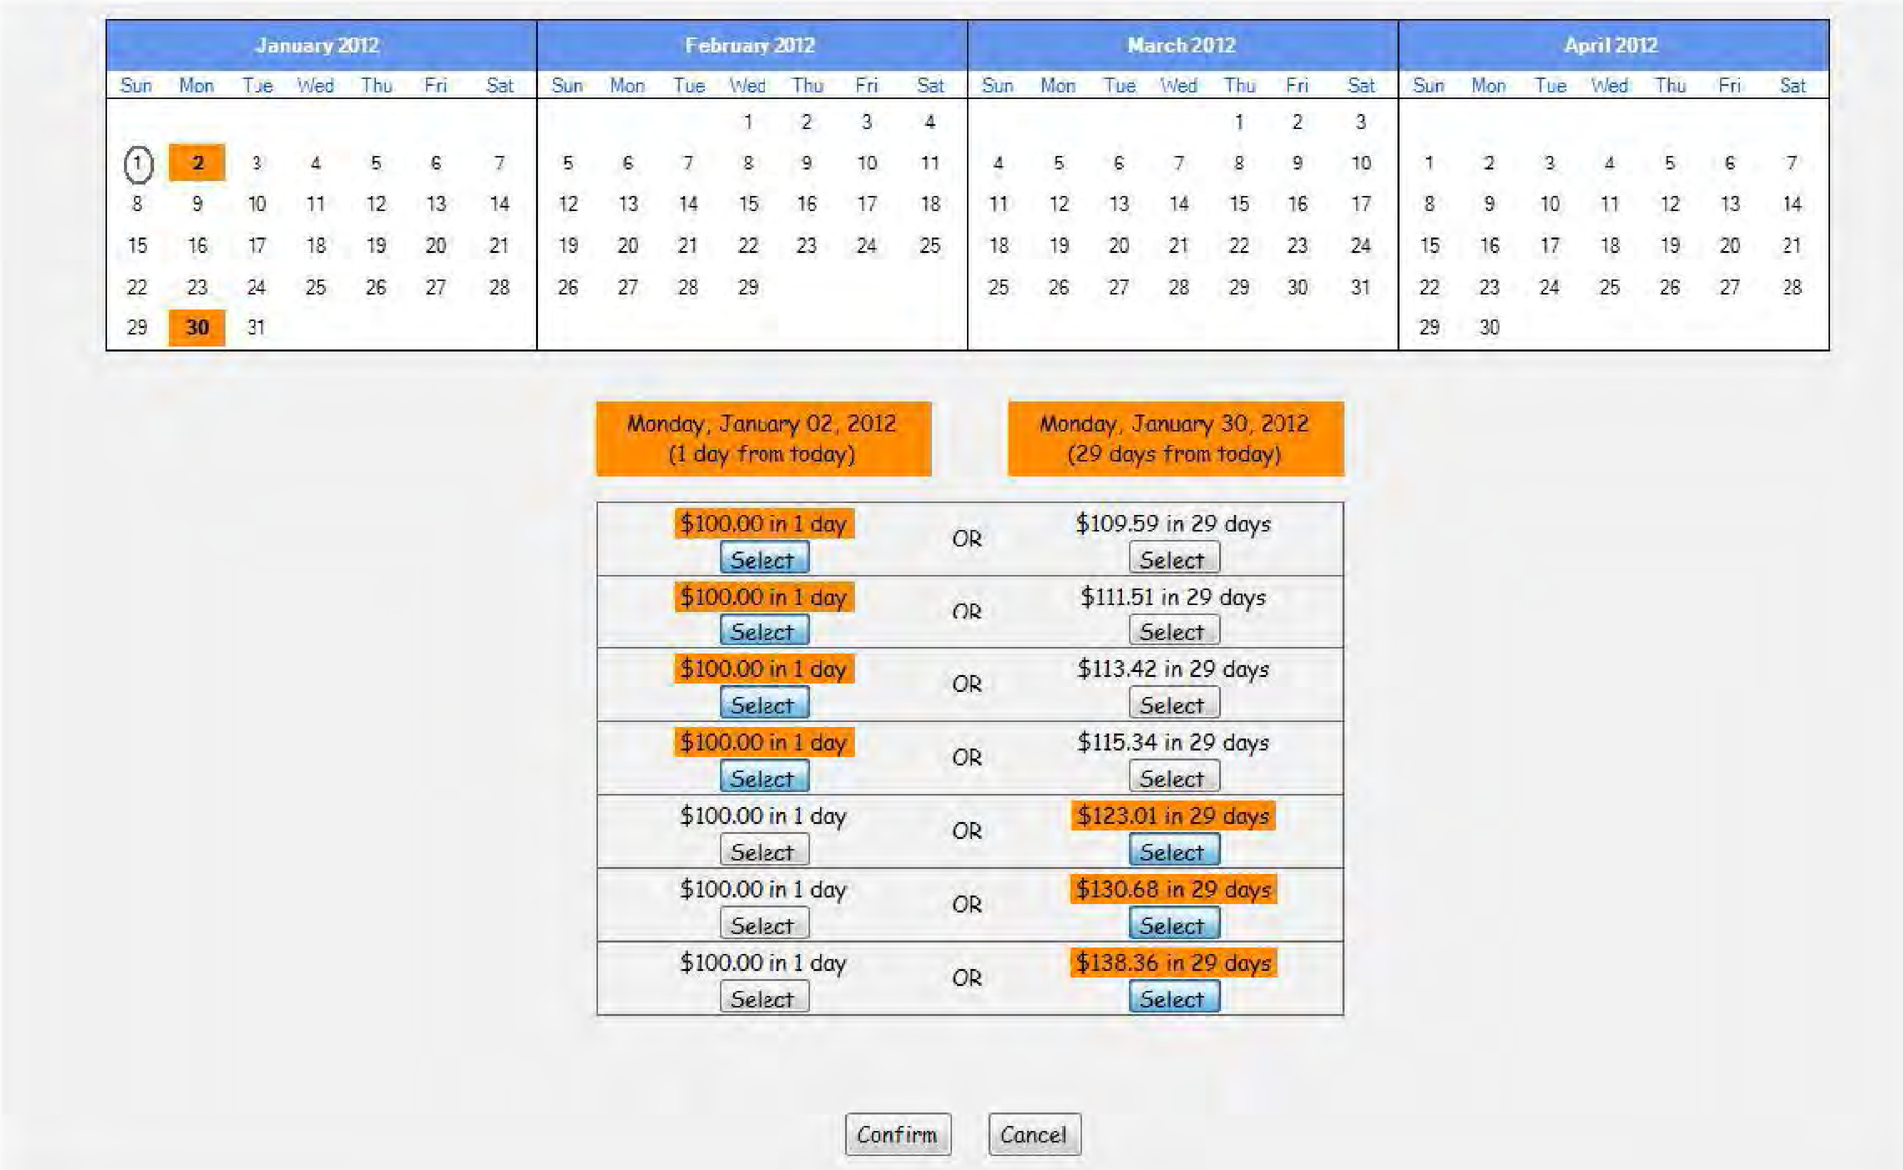
January 2012
February 2012
March 2012
April 2012
Sun Mon Tue Wed Thu Fri Sat
Monday, January 02, 2012 (1 day from today)
Monday, January 30, 2012 (29 days from today)
$100.00 in 1 day
Select
OR
$109.59 in 29 days
Select
$100.00 in 1 day
Select
OR
$111.51 in 29 days
Select
$100.00 in 1 day
Select
OR
$113.42 in 29 days
Select
$100.00 in 1 day
Select
OR
$115.34 in 29 days
Select
$100.00 in 1 day
Select
OR
$123.01 in 29 days
Select
$100.00 in 1 day
Select
OR
$130.68 in 29 days
Select
$100.00 in 1 day
Select
OR
$138.36 in 29 days
Select
Confirm
Cancel

This screen shows seven decisions. Each decision is presented on a different row. All decisions have the same format. For the purpose of explaining this task, assume for the moment that today is January 1, 2012. In the calendar, today's date is enclosed in a black circle. Let's look at the first decision in the example above (the one on the first decision row). The sooner choice pays $100 tomorrow, January 2, 2012 in the example, and the later choice pays $109.59 in twenty-nine days from today. You choose between these two options by clicking the button under the option you prefer.

We will present you with ten of these decision screens, with each screen having seven choices for you to make. You must make all seven choices on each decision screen before moving to the next decision screen. While on a single decision screen, the only difference between decisions is that the dollar amounts of the future payment will change. However, different decision screens will have different dollar amounts and future payment dates. So, you should make sure to pay attention to both the changing dollar amounts and changing dates as you make your decisions.

After you have worked through all of the decisions, we will select one of your ten decision screens by rolling a 10-sided die. Then we will roll a 10-sided die again, until a number between 1 and 7 comes up, to pick one decision row on that screen. When you make your choices you will not know which decision will be selected for payment. You should therefore treat each decision as if it might actually count for payment.

Once the decision screen and row are selected, we will look at the specific choice you made and the payout for that choice. If the payout is smaller (less than $100), you will actually be paid this amount. However, if the payout is larger ($100 or greater), you will roll a 10-sided die to determine whether or not you are actually paid this amount. If you roll a 1, you will actually be paid the amount, and on the date that you chose to receive it. If you roll a number other than 1, you will earn nothing in this task. This roll will be at the end of the experiment, and in private, when you are being paid.

For instance, suppose the decision screen in the above example was selected and you preferred the sooner date in the first four rows and the later date in the last three rows, as shown above. You would then roll a 10-sided die, until a number between 1 and 7 comes up, to select the row. Suppose the outcome was 5 and the fifth row is selected, where you preferred a payment of $123.01 in twenty-nine days. Since this payment is $100 or greater, you would roll the 10-sided die again, at the end of the experiment. If the outcome of your die-roll is 1, you will be paid $123.01 in twenty-nine days. However, if the outcome is 3, or anything other than 1, you will get nothing.

You will receive the money on the date stated in your preferred option. We will pay you using *PayPal*, which is an online payment service. We will explain more about *PayPal* in a few minutes. If you receive some money to be paid in the future you will also receive a written confirmation from Professor Harrison which guarantees that the money is to be paid to you on that date.

The money you receive from these choices is in addition to the show-up fee of $5, which is paid out at the end of the experiment as cash.

C. Instructions: Treatment T2

B_T2

Sooner versus Later Payments

In this task you will make a number of choices between receiving an amount of money on a "sooner" date or a different amount of money on a "later" date. The sooner date will always be today, while the later date will be some weeks from today. An example of a decision screen is shown below. The dates and amounts in your task will differ. You will make all decisions on a computer.

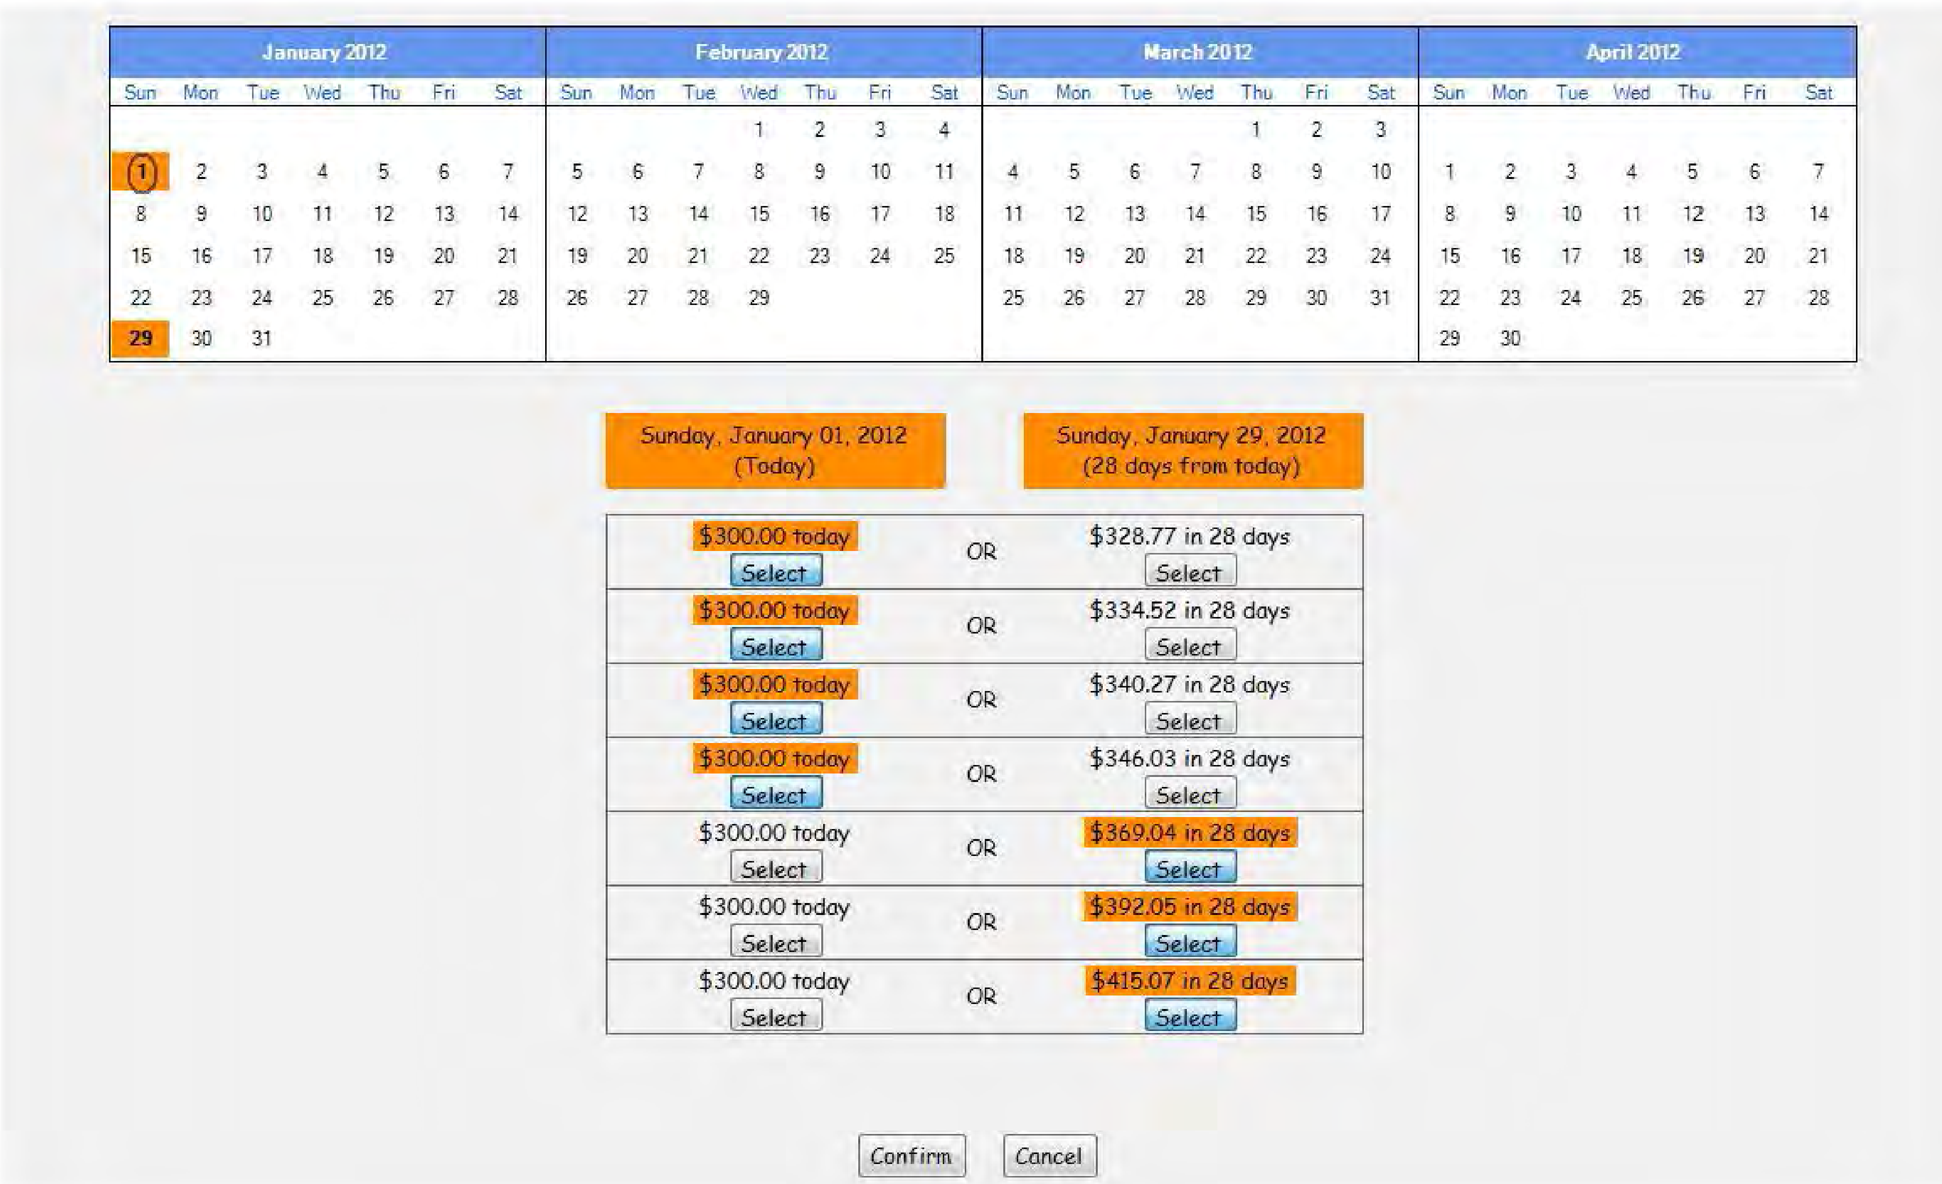
January 2012
February 2012
March 2012
April 2012
Sun Mon Tue Wed Thu Fri Sat
Sunday, January 01, 2012 (Today)
Sunday, January 29, 2012 (28 days from today)
$300.00 today OR $328.77 in 28 days
$300.00 today OR $334.52 in 28 days
$300.00 today OR $340.27 in 28 days
$300.00 today OR $346.03 in 28 days
$300.00 today OR $369.04 in 28 days
$300.00 today OR $392.05 in 28 days
$300.00 today OR $415.07 in 28 days
Select
Confirm
Cancel

This screen shows seven decisions. Each decision is presented on a different row. All decisions have the same format. For the purpose of explaining this task, assume for the moment that today is January 1, 2012. In the calendar, today's date is enclosed in a black circle. Let's look at the first decision in the example above (the one on the first decision row). The sooner choice pays $300 today, January 1, 2012 in the example, and the later choice pays $328.77 in twenty-eight days from today. You choose between these two options by clicking the button under the option you prefer.

We will present you with ten of these decision screens, with each screen having seven choices for you to make. You must make all seven choices on each decision screen before moving to the next decision screen. While on a single decision screen, the only difference between decisions is that the dollar amounts of the future payment will change. However, different decision screens will have different dollar amounts and future payment dates. So, you should make sure to pay attention to both the changing dollar amounts and changing dates as you make your decisions.

After you have worked through all of the decisions, we will select one of your ten decision screens by rolling a 10-sided die. Then we will roll a 10-sided die again, until a number between 1 and 7 comes up, to pick one decision row on that screen. When you make your choices you will not know which decision will be selected for payment. You should therefore treat each decision as if it might actually count for payment.

Once the decision screen and row are selected, we will look at the specific choice you made and the payout for that choice. If the payout is smaller (less than $300), you will actually be paid this amount. However, if the payout is larger ($300 or greater), you will roll a 10-sided die to determine whether or not you are actually paid this amount. If you roll a 1, you will actually be paid the amount, and on the date that you chose to receive it. If you roll a number other than 1, you will earn nothing in this task. This roll will be at the end of the experiment, and in private, when you are being paid.

For instance, suppose the decision screen in the above example was selected and you preferred the sooner date in the first four rows and the later date in the last three rows, as shown above. You would then roll a 10-sided die, until a number between 1 and 7 comes up, to select the row. Suppose the outcome was 5 and the fifth row is selected, where you preferred a payment of $369.04 in twenty-eight days. Since this payment is $300 or greater, you would roll the 10-sided die again, at the end of the experiment. If the outcome of your die-roll is 1, you will be paid $369.04 in twenty-eight days. However, if the outcome is 3, or anything other than 1, you will get nothing.

You will receive the money on the date stated in your preferred option. We will pay you using *PayPal*, which is an online payment service. We will explain more about *PayPal* in a few minutes. If you receive some money to be paid in the future you will also receive a written confirmation from Professor Harrison which guarantees that the money is to be paid to you on that date.

The money you receive from these choices is in addition to the show-up fee of $5, which is paid out at the end of the experiment as cash.

D. Instructions: Treatment T3

B_T3

Sooner versus Later Payments

In this task you will make a number of choices between receiving an amount of money on a "sooner" date or a different amount of money on a "later" date. The sooner date will be some weeks from today, while the later date will be even more weeks from today. An example of a decision screen is shown below. The dates and amounts in your task will differ. You will make all decisions on a computer.

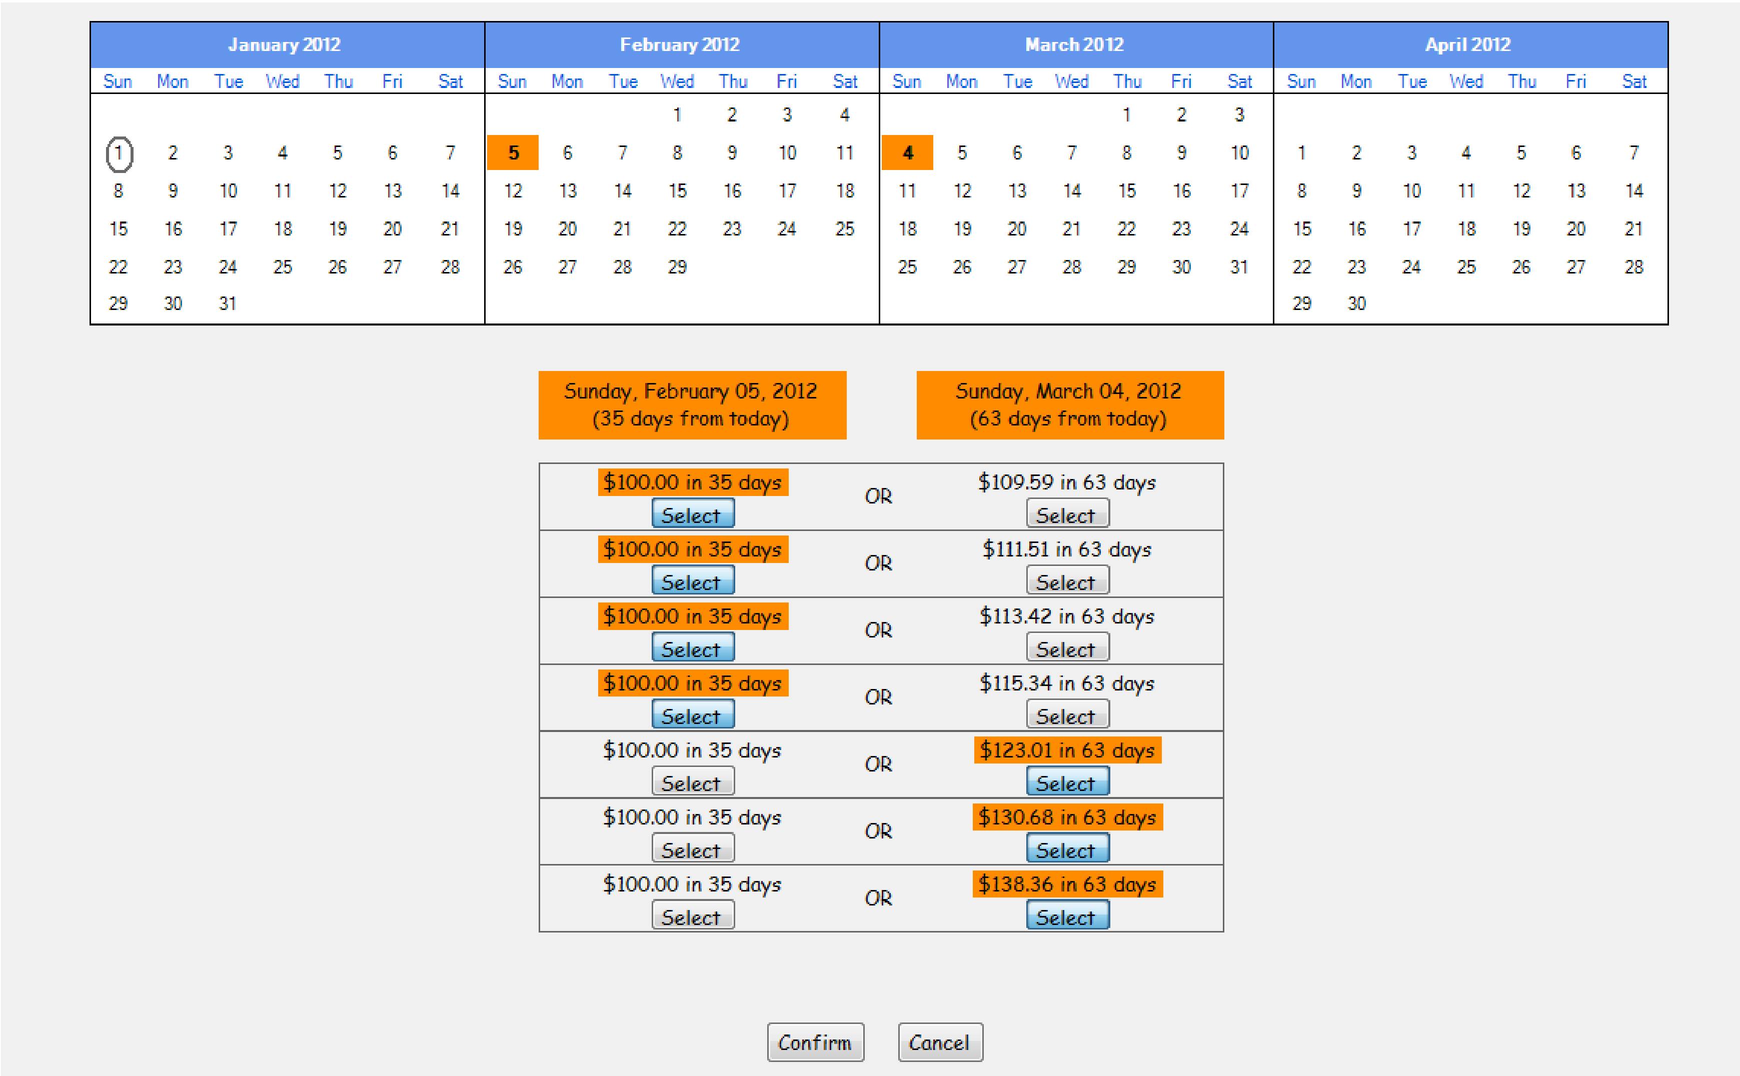
January 2012
February 2012
March 2012
April 2012
Sun Mon Tue Wed Thu Fri Sat
Sunday, February 05, 2012
(35 days from today)
Sunday, March 04, 2012
(63 days from today)
$100.00 in 35 days OR $109.59 in 63 days
$100.00 in 35 days OR $111.51 in 63 days
$100.00 in 35 days OR $113.42 in 63 days
$100.00 in 35 days OR $115.34 in 63 days
$100.00 in 35 days OR $123.01 in 63 days
$100.00 in 35 days OR $130.68 in 63 days
$100.00 in 35 days OR $138.36 in 63 days
Select
Confirm
Cancel

This screen shows seven decisions. Each decision is presented on a different row. All decisions have the same format. For the purpose of explaining this task, assume for the moment that today is January 1, 2012. In the calendar, today's date is enclosed in a black circle. Let's look at the first decision in the example above (the one on the first decision row). The sooner choice pays $100 in thirty-five days from today, where "today" is January 1, 2012 in the example, and the later choice pays $109.59 in sixty-three days from today. You choose between these two options by clicking the button under the option you prefer.

We will present you with ten of these decision screens, with each screen having seven choices for you to make. You must make all seven choices on each decision screen before moving to the next decision screen. While on a single decision screen, the only difference between decisions is that the dollar amounts of the future payment will change. However, different decision screens will have different dollar amounts and future payment dates. So, you should make sure to pay attention to both the changing dollar amounts and changing dates as you make your decisions.

After you have worked through all of the decisions, we will select one of your ten decision screens by rolling a 10-sided die. Then we will roll a 10-sided die again, until a number between 1 and 7 comes up, to pick one decision row on that screen. When you make your choices you will not know which decision will be selected for payment. You should therefore treat each decision as if it might actually count for payment.

Once the decision screen and row are selected, we will look at the specific choice you made and the payout for that choice. If the payout is smaller (less than $100), you will actually be paid this amount. However, if the payout is larger ($100 or greater), you will roll a 10-sided die to determine whether or not you are actually paid this amount. If you roll a 1, you will actually be paid the amount, and on the date that you chose to receive it. If you roll a number other than 1, you will earn nothing in this task. This roll will be at the end of the experiment, and in private, when you are being paid.

For instance, suppose the decision screen in the above example was selected and you preferred the sooner date in the first four rows and the later date in the last three rows, as shown above. You would then roll a 10-sided die, until a number between 1 and 7 comes up, to select the row. Suppose the outcome was 5 and the fifth row is selected, where you preferred a payment of $123.01 in sixty-three days. Since this payment is $100 or greater, you would roll the 10-sided die again, at the end of the experiment. If the outcome of your die-roll is 1, you will be paid $123.01 in sixty-three days. However, if the outcome is 3, or anything other than 1, you will get nothing.

You will receive the money on the date stated in your preferred option. We will pay you using *PayPal*, which is an online payment service. We will explain more about *PayPal* in a few minutes. If you receive some money to be paid in the future you will also receive a written confirmation from Professor Harrison which guarantees that the money is to be paid to you on that date.

The money you receive from these choices is in addition to the show-up fee of $5, which is paid out at the end of the experiment as cash.

E. Instructions: Treatment T4

B_T4

Sooner versus Later Payments

In this task you will make a number of choices between receiving an amount of money on a "sooner" date or a different amount of money on a "later" date. You will be presented with a series of decision pairs arranged in seven rows per decision screen. Each decision consists of a smaller amount that will be paid sooner and a larger amount that will be paid at a later date. The two pairs in a row involve the same amounts but different dates. An example of a decision screen is shown below. The dates and amounts in your task will differ. You will make all decisions on a computer.

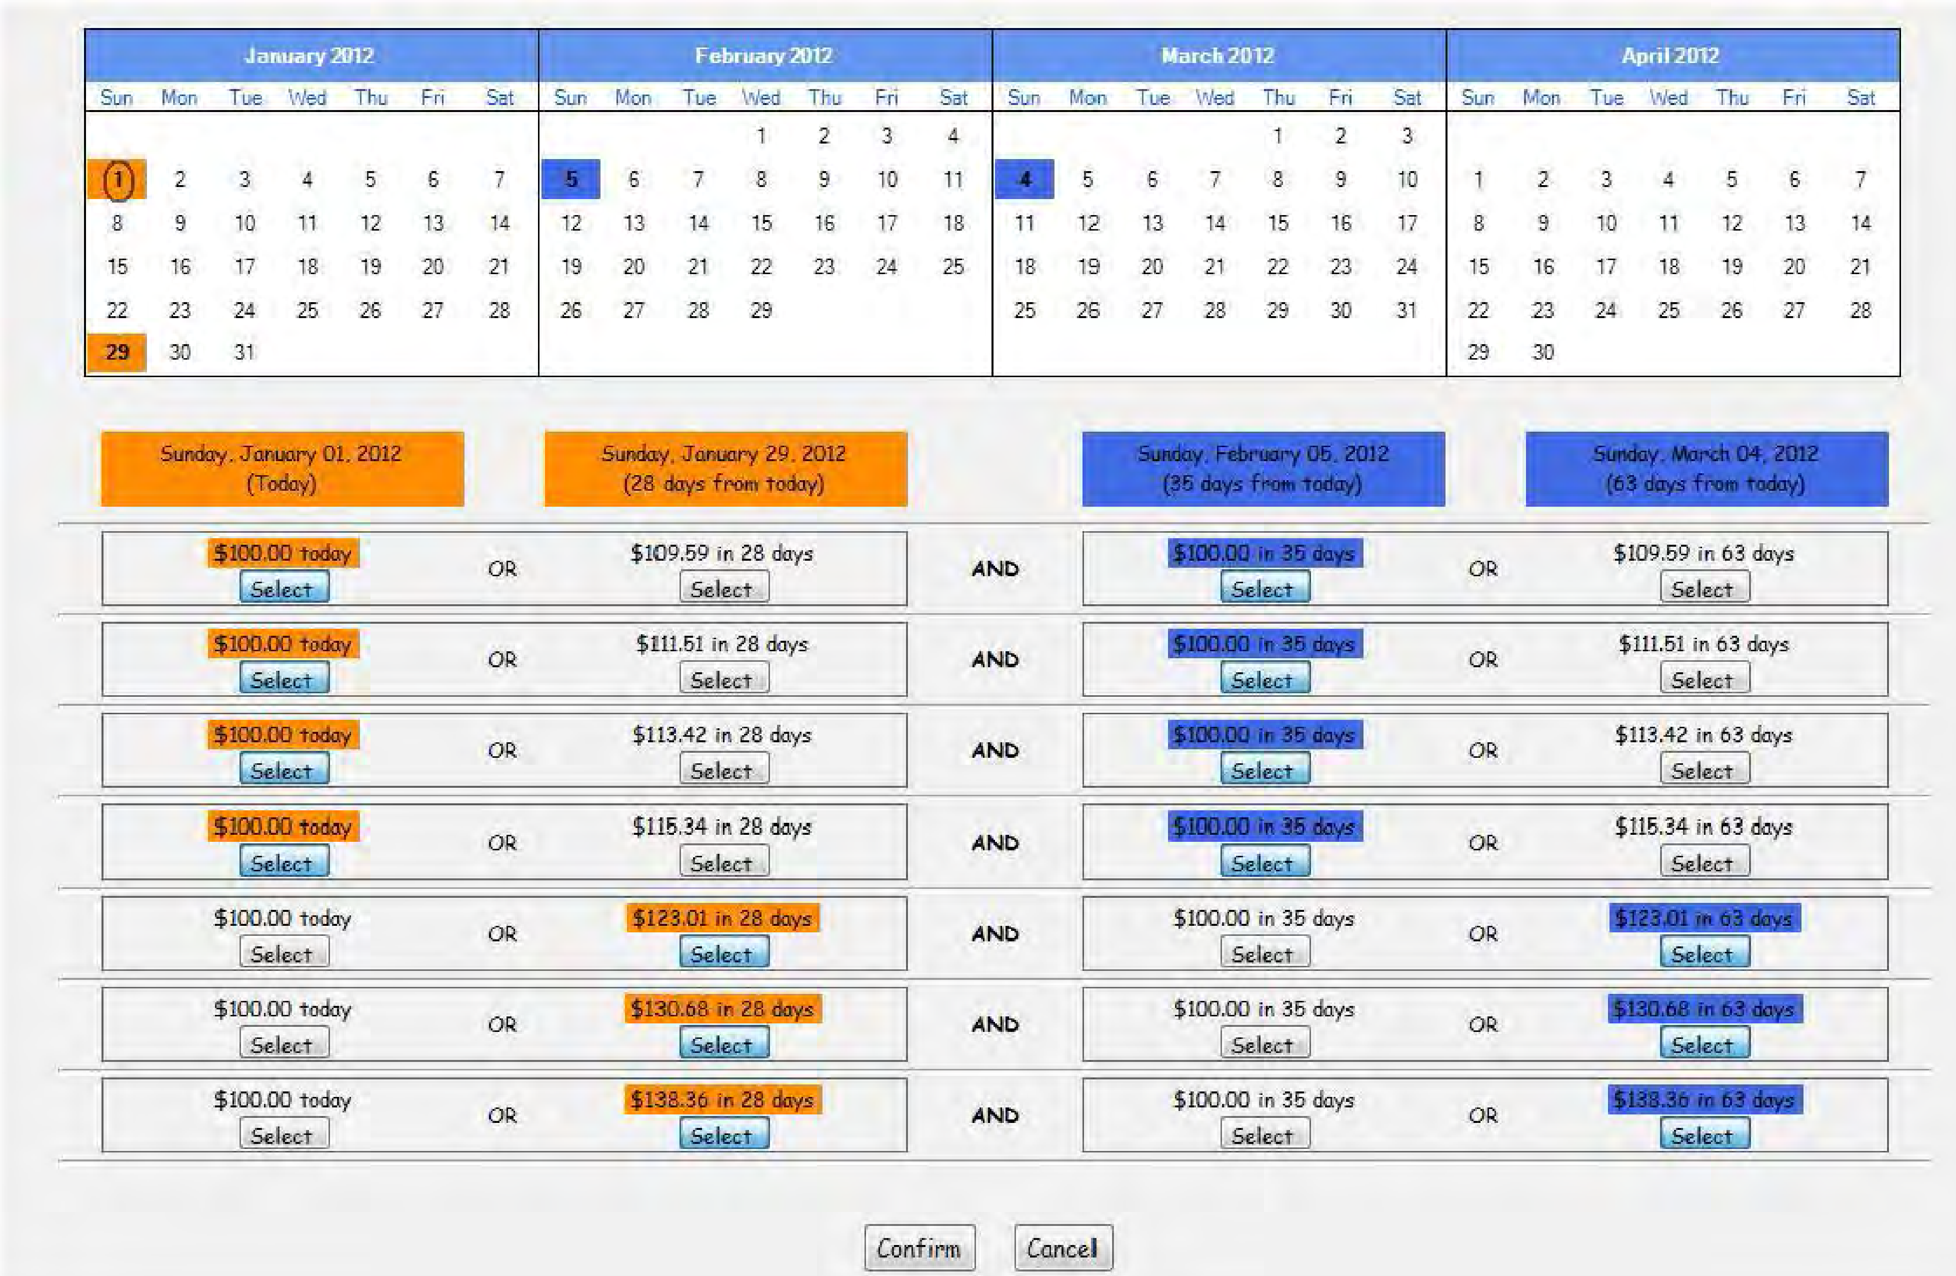
January 2012
February 2012
March 2012
April 2012
Sunday, January 01, 2012 (Today)
Sunday, January 29, 2012 (28 days from today)
Sunday, February 05, 2012 (35 days from today)
Sunday, March 04, 2012 (63 days from today)
$100.00 today OR $109.59 in 28 days AND $100.00 in 35 days OR $109.59 in 63 days
$100.00 today OR $111.51 in 28 days AND $100.00 in 35 days OR $111.51 in 63 days
$100.00 today OR $113.42 in 28 days AND $100.00 in 35 days OR $113.42 in 63 days
$100.00 today OR $115.34 in 28 days AND $100.00 in 35 days OR $115.34 in 63 days
$100.00 today OR $123.01 in 28 days AND $100.00 in 35 days OR $123.01 in 63 days
$100.00 today OR $130.68 in 28 days AND $100.00 in 35 days OR $130.68 in 63 days
$100.00 today OR $138.36 in 28 days AND $100.00 in 35 days OR $138.36 in 63 days
Select
Confirm
Cancel

This screen shows seven pairs of independent decisions. Each decision pair is presented on a different row. All decision pairs have the same format. For the purpose of explaining this task, assume for the moment that today is January 1, 2012. In the calendar, today's date is enclosed in a black circle. Let's look at the first decision pair in the example above (the one on the first decision row). In the first decision, the sooner choice pays $100 today, January 1, 2012 in the example, and the later choice pays $109.59 in twenty-eight days from today. The second decision in this row offers the choice between $100 in thirty-five days or $109.59 in sixty-three days from today.

If you choose the sooner option for one decision in a pair, you are free to choose the sooner or later option for the other decision in that pair. And if you choose the later option for one decision in a pair, you are free to choose the sooner or later option for the other decision in that pair. In other words, it is possible to choose the sooner option in one decision and the later option in the other decision of the same pair, to choose the later option in one decision and the sooner option the other decision of the same pair, or to choose the same option in each decision of the same pair. You choose by clicking the button under the alternative you prefer.

The above screen example shows a situation in which someone chose the same options in each pair. Here is a screen example in which someone chose differently for the 3rd and 4th decision rows:

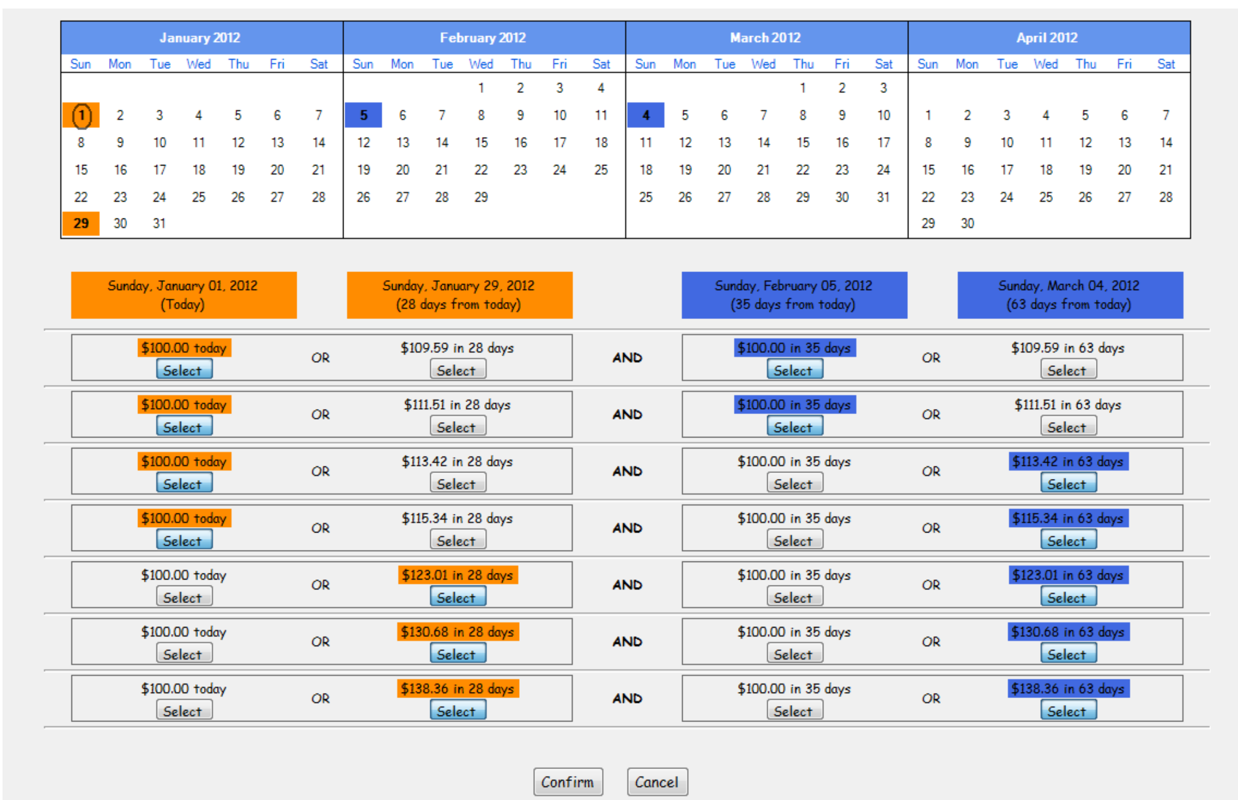

And here is an example in which someone chose differently for the 5th and 6th decision rows:

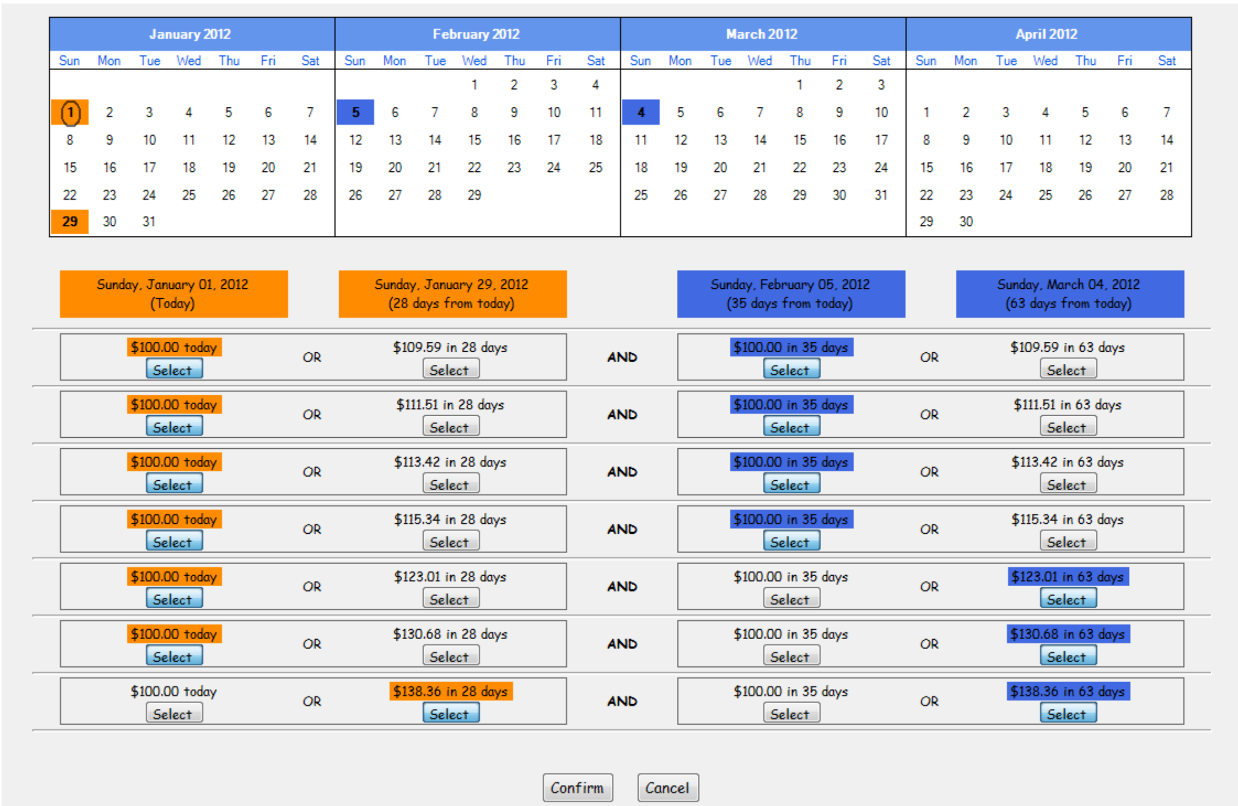

We will present you with ten of these decision screens, with each screen displaying seven decision pairs and fourteen choices for you to make. You must make all fourteen choices on each decision screen before moving to the next decision screen. While on a single decision screen, the only difference between decision pairs is that the dollar amounts of the future payment will change. However, pairs on different decision screens will have different dollar amounts and future payment dates. So, you should make sure to pay attention to both the changing dollar amounts and changing dates as you make your decisions.

After you have worked through all of the decisions, we will select one of your ten decision screens by rolling a 10-sided die. Then we will roll a 10-sided die again, until a number between 1 and 7 comes up, to pick one decision row on that screen. This means you will be paid for both choices that you made in that pair. When you make your choices you will not know which decision pair will be selected for payment. You should therefore treat each decision as if it might actually count for payment.

Once the decision screen and row are selected, we will look at the specific choices you made and the payouts for those choices. If the individual payouts in the selected decision row are smaller (less than $100), you will actually be paid both amounts. However, if the individual payouts in the selected decision row are larger ($100 or greater), you will roll a 10-sided die to determine whether or not you are actually paid these amounts. If you roll a 1, you will actually be paid the amounts, and on the dates that you chose to receive them. If you roll a number other than 1, you will earn nothing in this task. This roll will be at the end of the experiment, and in private, when you are being paid.

For instance, suppose the decision screen in the first screen example was selected and you preferred the sooner date in the first four rows and the later date in the last three rows, as shown above. You would then roll a 10-sided die, until a number between 1 and 7 comes up, to select the row. Suppose the outcome was 5 and the fifth row is selected. In this row you preferred a payment of $123.01 in twenty-eight days in the first decision and $123.01 in sixty-three days in the second decision. Since this payment is $100 or greater, you would roll the 10-sided die again, at the end of the experiment. If the outcome of your die-roll is 1, you will be paid $123.01 in twenty-eight days and another $123.01 in sixty-three days. However, if the outcome is 3, or anything other than 1, you will get nothing.

You will receive the money on the date stated in your preferred option. We will pay you using *PayPal*, which is an online payment service. We will explain more about *PayPal* in a few minutes. If you receive some money to be paid in the future you will also receive a written confirmation from Professor Harrison which guarantees that the money is to be paid to you on that date.

The money you receive from these choices is in addition to the show-up fee of $5, which is paid out at the end of the experiment as cash.

F. Instructions: Treatment T5

B_T5

Sooner versus Later Payments

In this task you will make a number of choices between receiving an amount of money on a "sooner" date or a different amount of money on a "later" date. You will be presented with a series of decision pairs arranged in seven rows per decision screen. Each decision consists of a smaller amount that will be paid sooner and a larger amount that will be paid at a later date. The two pairs in a row involve the same amounts but different dates. An example of a decision screen is shown below. The dates and amounts in your task will differ. You will make all decisions on a computer.

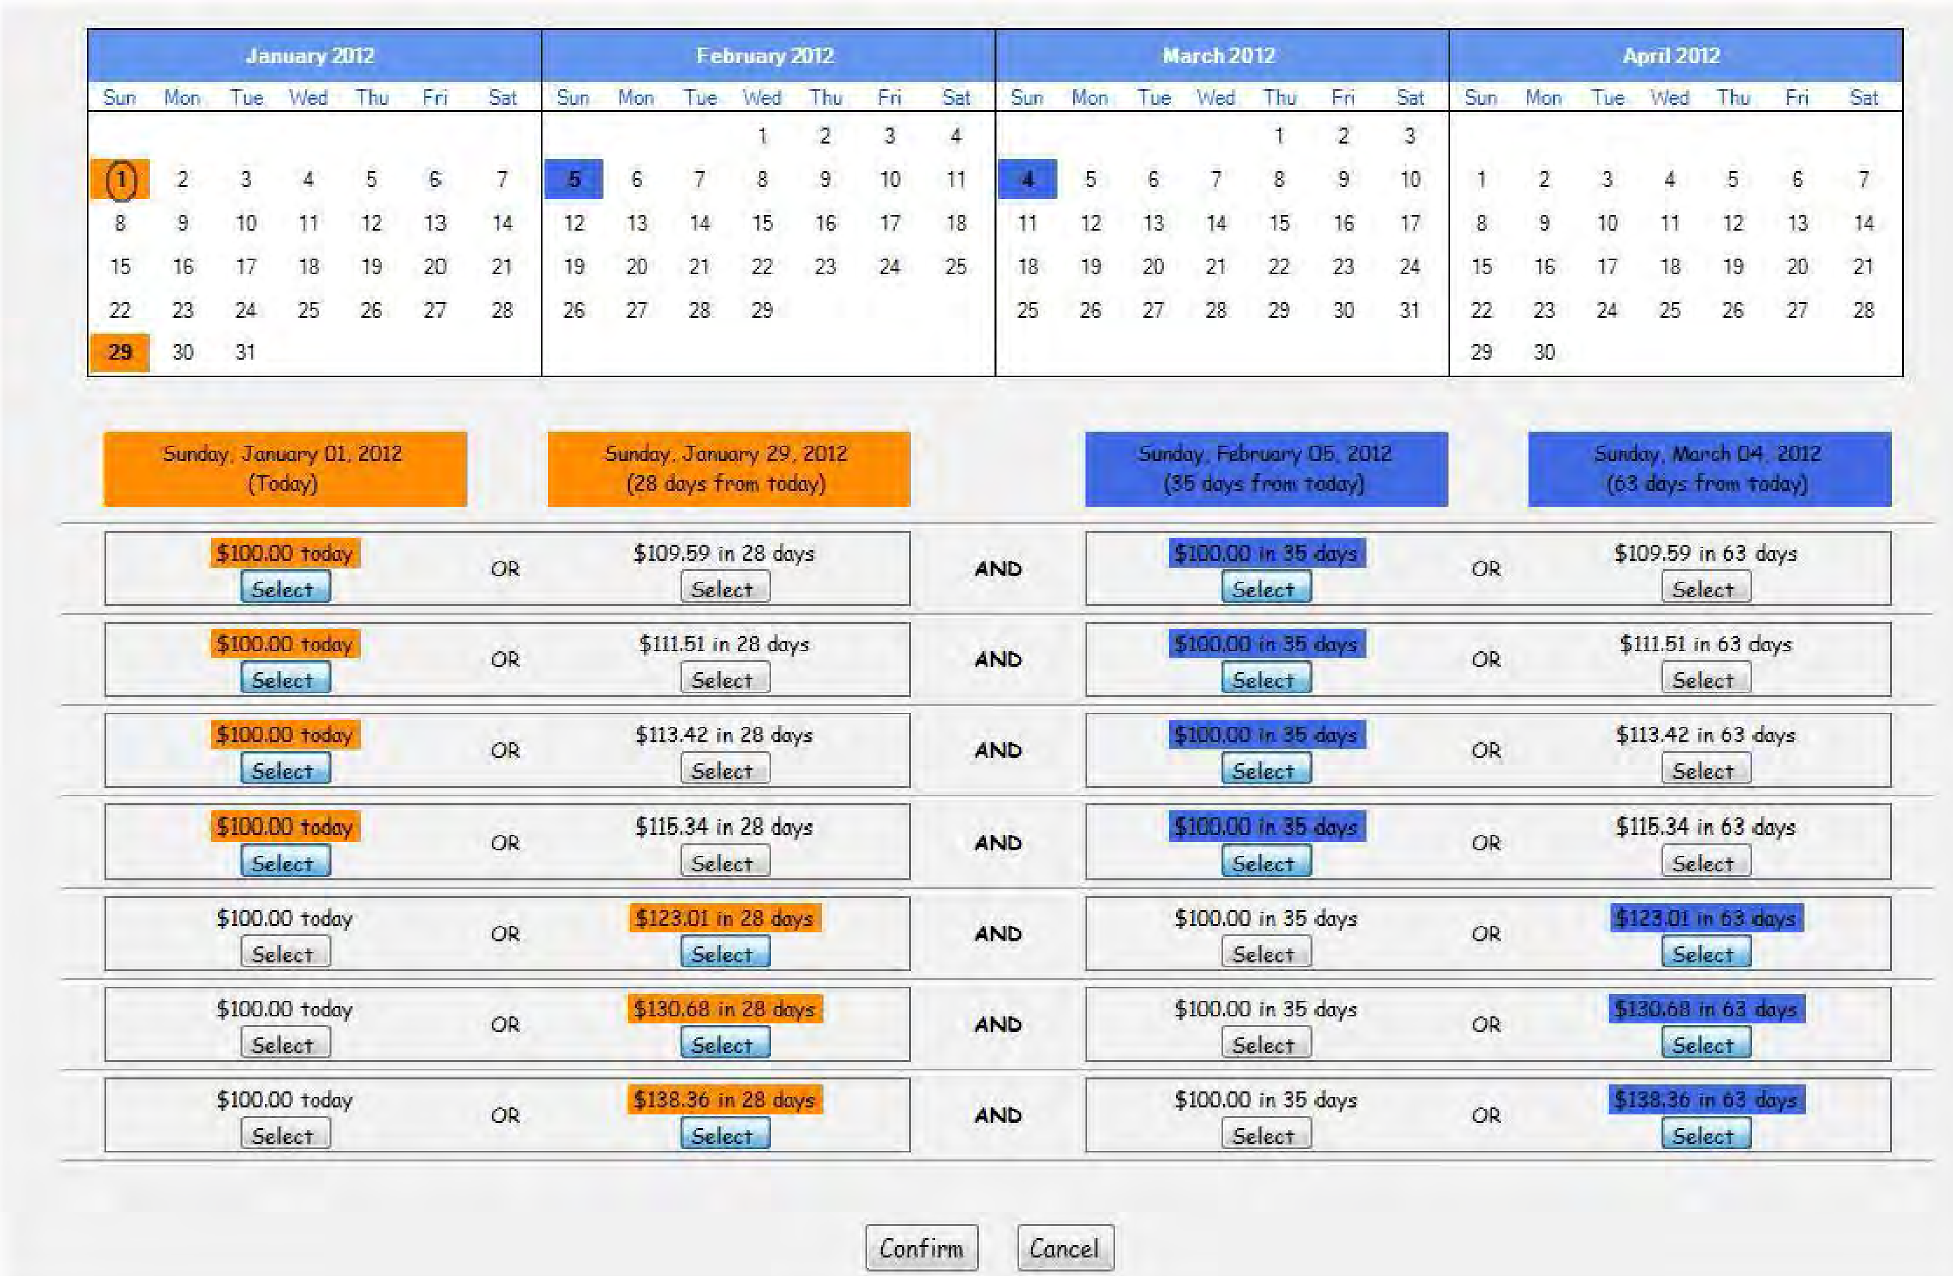
January 2012
February 2012
March 2012
April 2012
Sun Mon Tue Wed Thu Fri Sat
Sunday, January 01, 2012 (Today)
Sunday, January 29, 2012 (28 days from today)
Sunday, February 05, 2012 (35 days from today)
Sunday, March 04, 2012 (63 days from today)
$100.00 today OR $109.59 in 28 days AND $100.00 in 35 days OR $109.59 in 63 days
$100.00 today OR $111.51 in 28 days AND $100.00 in 35 days OR $111.51 in 63 days
$100.00 today OR $113.42 in 28 days AND $100.00 in 35 days OR $113.42 in 63 days
$100.00 today OR $115.34 in 28 days AND $100.00 in 35 days OR $115.34 in 63 days
$100.00 today OR $123.01 in 28 days AND $100.00 in 35 days OR $123.01 in 63 days
$100.00 today OR $130.68 in 28 days AND $100.00 in 35 days OR $130.68 in 63 days
$100.00 today OR $138.36 in 28 days AND $100.00 in 35 days OR $138.36 in 63 days
Select
Confirm
Cancel

This screen shows seven pairs of tied-together decisions. Each decision pair is presented on a different row. All decision pairs have the same format. For the purpose of explaining this task, assume for the moment that today is January 1, 2012. In the calendar, today's date is enclosed in a black circle. Let's look at the first decision pair in the example above (the one on the first decision row). The first sooner choice pays $100 today, January 1, 2012 in the example, and the first later choice pays $109.59 in twenty-eight days from today. The second sooner choice pays $100 in thirty-five days from today and the second later choice pays $109.59 in sixty-three days from today.

If you choose the sooner option for one decision in a pair, you are also choosing the sooner option for the other decision in that pair. In other words, it is not possible to choose the sooner option in one decision and the later option in the other decision of the same pair. If you select the sooner options in one pair, you are free to select the later options in any other pair. You choose by clicking the button under the alternative you prefer.

We will present you with ten of these decision screens, with each screen displaying seven choices for you to make. You must make all seven choices on each decision screen before moving to the next decision screen. While on a single decision screen, the only difference between decision pairs is that the dollar amounts of the future payment will change. However, pairs on different decision screens will have different dollar amounts and future payment dates. So, you should make sure to pay attention to both the changing dollar amounts and changing dates as you make your decisions.

After you have worked through all of the decisions, we will select one of your ten decision screens by rolling a 10-sided die. Then we will roll a 10-sided die again, until a number between 1 and 7 comes up, to pick one decision row on that screen. This means you will be paid for both of the choices in that pair, which you made at the same time. When you make your choices you will not know which decision pair will be selected for payment. You should therefore treat each decision as if it might actually count for payment.

Once the decision screen and row are selected, we will look at the specific choice you made and the resulting payouts. If the individual payouts in the selected decision row are smaller (less than $100), you will actually be paid both amounts. However, if the individual payouts in the selected decision row are larger ($100 or greater), you will roll a 10-sided die to determine whether or not you are actually paid these amounts. If you roll a 1, you will actually be paid the amounts, and on the dates that you chose to receive them. If you roll a number other than 1, you will earn nothing in this task. This roll will be at the end of the experiment, and in private, when you are being paid.

For instance, suppose the decision screen in the above example was selected and you preferred the sooner date in the first four rows and the later date in the last three rows, as shown above. You would then roll a 10-sided die, until a number between 1 and 7 comes up, to select the row. Suppose the outcome was 5 and the fifth row is selected. In this row you preferred a payment of $123.01 in twenty-eight days in the first decision and, consequently, $123.01 in sixty-three days in the second decision. Since this payment is $100 or greater, you would roll the 10-sided die again, at the end of the experiment. If the outcome of your die-roll is 1, you will be paid $123.01 in twenty-eight days and another $123.01 in sixty-three days. However, if the outcome is 3, or anything other than 1, you will get nothing.

You will receive the money on the date stated in your preferred option. We will pay you using *PayPal*, which is an online payment service. We will explain more about *PayPal* in a few minutes. If you receive some money to be paid in the future you will also receive a written confirmation from Professor Harrison which guarantees that the money is to be paid to you on that date.

The money you receive from these choices is in addition to the show-up fee of $5, which is paid out at the end of the experiment as cash.

G. Instructions: Treatment T6

B_T6

Sooner versus Later Payments

In this task you will make a number of choices between receiving an amount of money on a "sooner" date or a different amount of money on a "later" date. You will be presented with a series of decision triples arranged in seven rows per decision screen. Each decision consists of a smaller amount that will be paid sooner and a larger amount that will be paid at a later date. The three triples in a row involve the same amounts but different dates. An example of a decision screen is shown below. The dates and amounts in your task will differ. You will make all decisions on a computer.

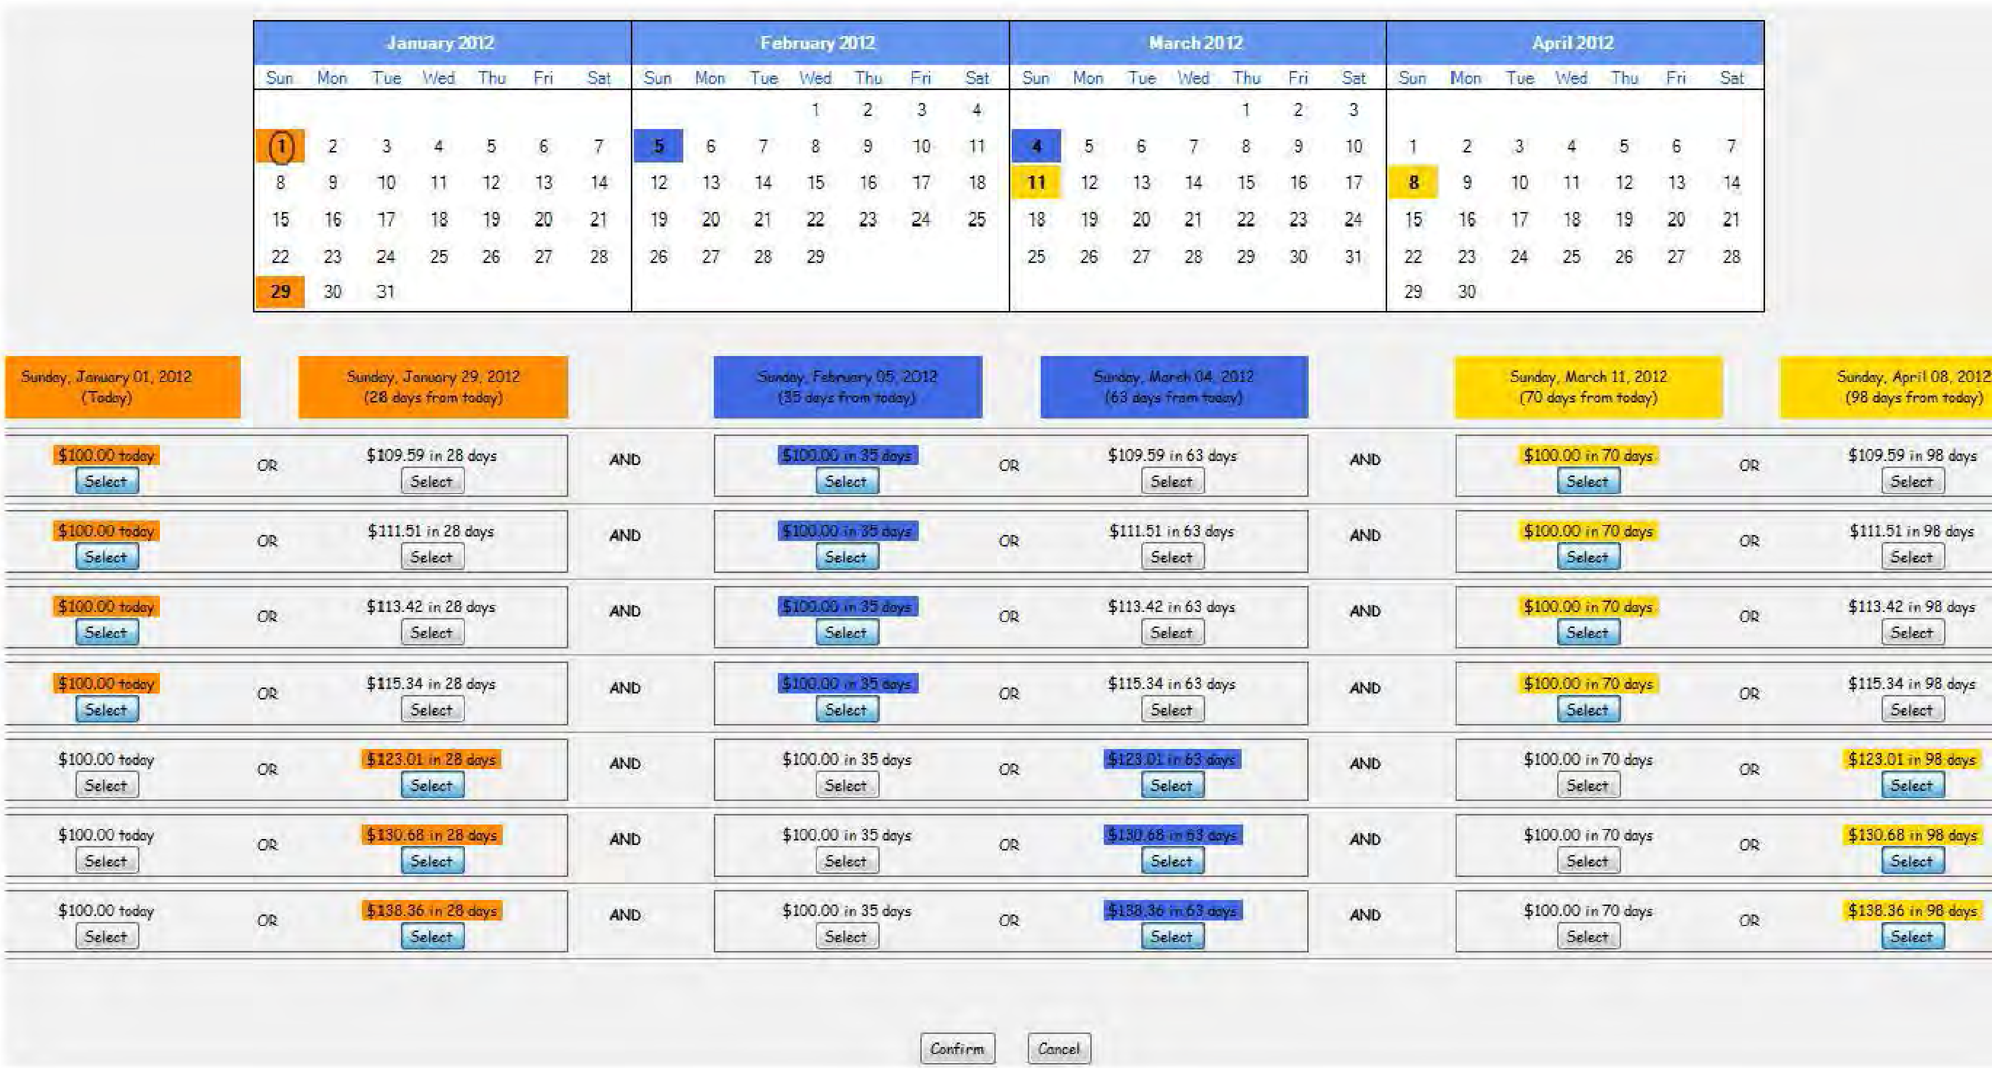
January 2012
February 2012
March 2012
April 2012
Sunday, January 01, 2012 (Today)
Sunday, January 29, 2012 (28 days from today)
Sunday, February 05, 2012 (35 days from today)
Sunday, March 04, 2012 (63 days from today)
Sunday, March 11, 2012 (70 days from today)
Sunday, April 08, 2012 (98 days from today)
$100.00 today OR $109.59 in 28 days AND $100.00 in 35 days OR $109.59 in 63 days AND $100.00 in 70 days OR $109.59 in 98 days
$100.00 today OR $111.51 in 28 days AND $100.00 in 35 days OR $111.51 in 63 days AND $100.00 in 70 days OR $111.51 in 98 days
$100.00 today OR $113.42 in 28 days AND $100.00 in 35 days OR $113.42 in 63 days AND $100.00 in 70 days OR $113.42 in 98 days
$100.00 today OR $115.34 in 28 days AND $100.00 in 35 days OR $115.34 in 63 days AND $100.00 in 70 days OR $115.34 in 98 days
$100.00 today OR $123.01 in 28 days AND $100.00 in 35 days OR $123.01 in 63 days AND $100.00 in 70 days OR $123.01 in 98 days
$100.00 today OR $130.68 in 28 days AND $100.00 in 35 days OR $130.68 in 63 days AND $100.00 in 70 days OR $130.68 in 98 days
$100.00 today OR $138.36 in 28 days AND $100.00 in 35 days OR $138.36 in 63 days AND $100.00 in 70 days OR $138.36 in 98 days
Select
Confirm
Cancel

This screen shows seven triples of independent decisions. Each decision triple is presented on a different row. All decision triples have the same format. For the purpose of explaining this task, assume for the moment that today is January 1, 2012. In the calendar, today's date is enclosed in a black circle. Let's look at the first decision triple in the example above (the one on the first decision row). In the first decision, the sooner choice pays $100 today, January 1, 2012 in the example, and the later choice pays $109.59 in twenty-eight days from today. The second decision in this row offers the choice between $100 in thirty-five days or $109.59 in sixty-three days from today. The third decision in this row offers the choice between $100 in seventy days or $109.59 in ninety-eight days from today.

We will present you with ten of these decision screens, with each screen displaying seven decision triples and twenty-one choices for you to make. You must make all twenty-one choices on each decision screen before moving to the next decision screen. While on a single decision screen, the only difference between decision triples is that the dollar amounts of the future payment will change. However, triples on different decision screens will have different dollar amounts and future payment dates. So, you should make sure to pay attention to both the changing dollar amounts and changing dates as you make your decisions.

If you choose the sooner option for one decision in a triple, you are free to choose the sooner or later option for the other two decisions in that triple. And if you choose the later option for one decision in a triple, you are free to choose the sooner or later option for the other two decisions in that triple. In other words, it is possible to choose the sooner option in one decision and the later option in any of the other two decisions of the same triple, to choose the later option in one decision and the sooner option in any of the other two decisions of the same triple, or to choose the same option in each decision of the same triple. You choose by clicking the button under the alternative you prefer.

The above screen example shows a situation in which someone chose the same options in each triple. Here is a screen example in which someone chose differently for the 3rd and 4th decision rows:

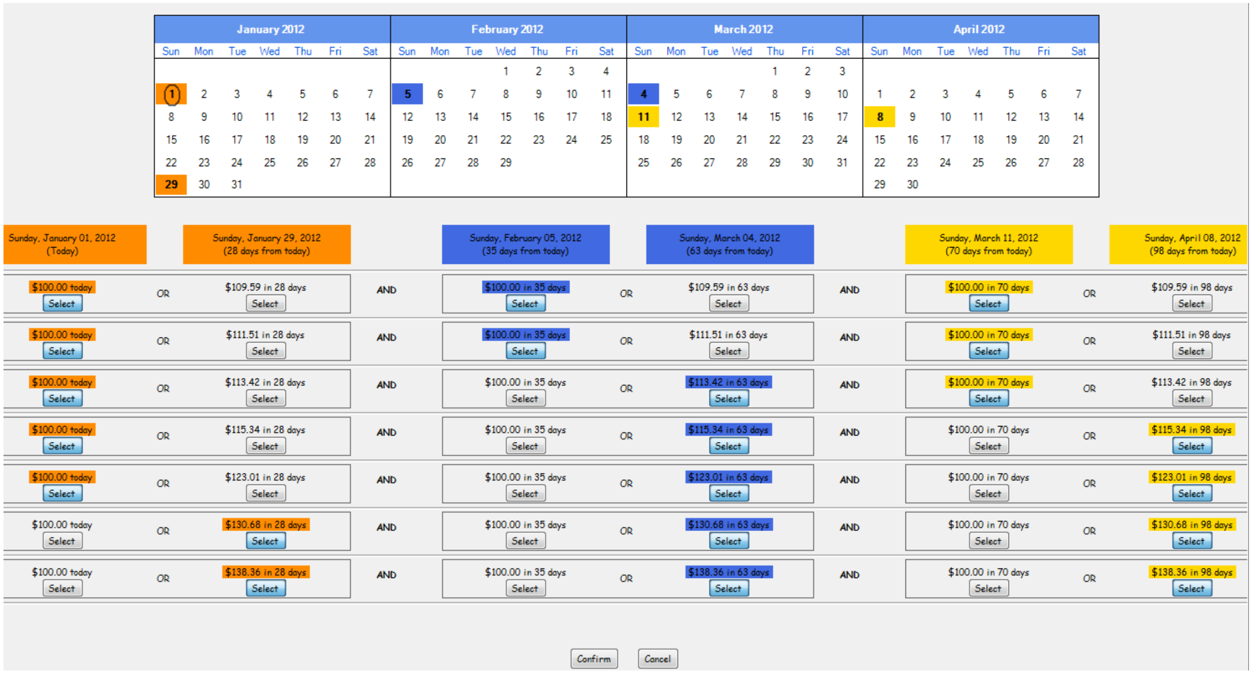

And here is an example in which someone chose differently for the 5th and 6th decision rows:

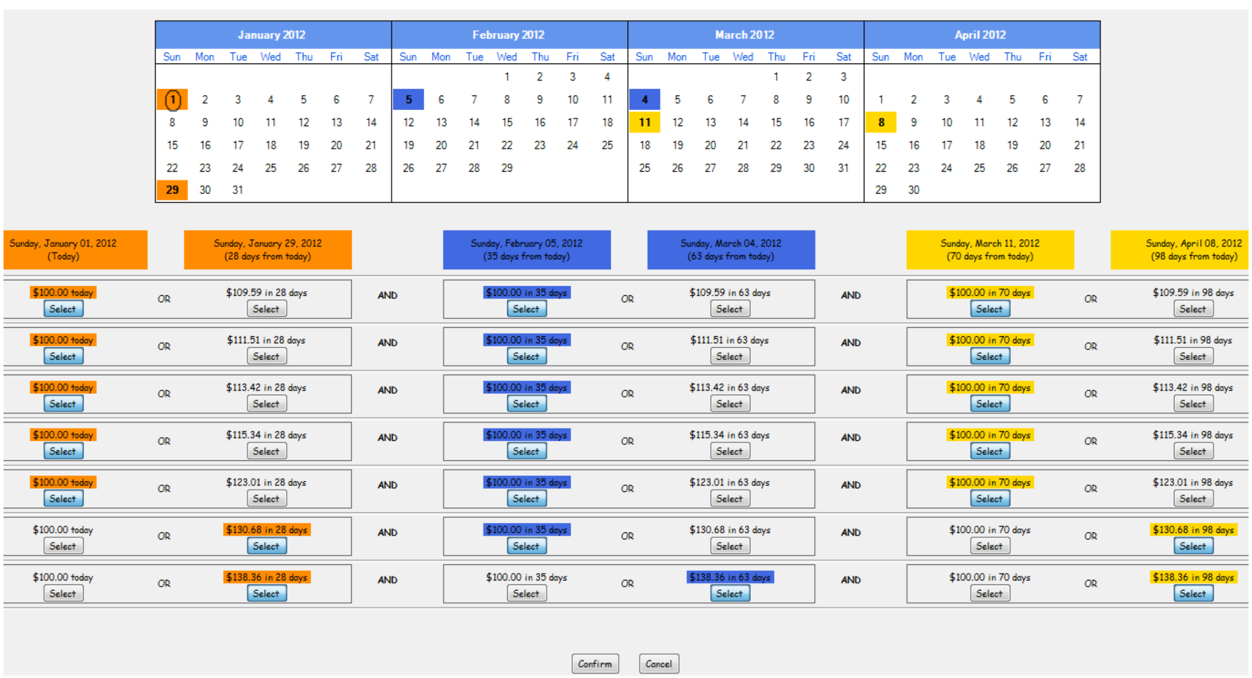

After you have worked through all of the decisions, we will select one of your ten decision screens by rolling a 10-sided die. Then we will roll a 10-sided die again, until a number between 1 and 7 comes up, to pick one decision row on that screen. This means you will be paid for <u>all three</u> choices that you made in that <u>triple</u>. When you make your choices you will not know which decision triple will be selected for payment. You should therefore treat each decision as if it might actually count for payment.

Once the decision screen and row are selected, we will look at the specific choices you made and the payouts for those choices. If the individual payouts in the selected decision row are smaller (less than $100), you will actually be paid all three amounts. However, if the individual payouts in the selected decision row are larger ($100 or greater), you will roll a 10-sided die to determine whether or not you are actually paid these amounts. If you roll a 1, you will actually be paid the amounts, and on the dates that you chose to receive them. If you roll a number other than 1, you will earn nothing in this task. This roll will be at the end of the experiment, and in private, when you are being paid.

For instance, suppose the decision screen in the first screen example was selected and you preferred the sooner date in the first four rows and the later date in the last three rows, as shown above. You would then roll a 10-sided die, until a number between 1 and 7 comes up, to select the row. Suppose the outcome was 5 and the fifth row is selected. In this row you preferred a payment of $123.01 in twenty-eight days in the first decision, $123.01 in sixty-three days in the second decision, and $123.01 in ninety-eight days in the third decision. Since this payment is $100 or greater, you would roll the 10-sided die again, at the end of the experiment. If the outcome of your die-roll is 1, you will be paid $123.01 in twenty-eight days, $123.01 in sixty-three days, and another $123.01. However, if the outcome is 3, or anything other than 1, you will get nothing.

You will receive the money on the date stated in your preferred option. We will pay you using *PayPal*, which is an online payment service. We will explain more about *PayPal* in a few minutes. If you receive some money to be paid in the future you will also receive a written confirmation from Professor Harrison which guarantees that the money is to be paid to you on that date.

The money you receive from these choices is in addition to the show-up fee of $5, which is paid out at the end of the experiment as cash.

H. *Instructions: Treatment T7*

B_T7

Sooner versus Later Payments

In this task you will make a number of choices between receiving an amount of money on a "sooner" date or a different amount of money on a "later" date. You will be presented with a series of decision triples arranged in seven rows per decision screen. Each decision consists of a smaller amount that will be paid sooner and a larger amount that will be paid at a later date. The three triples in a row involve the same amounts but different dates. An example of a decision screen is shown below. The dates and amounts in your task will differ. You will make all decisions on a computer.

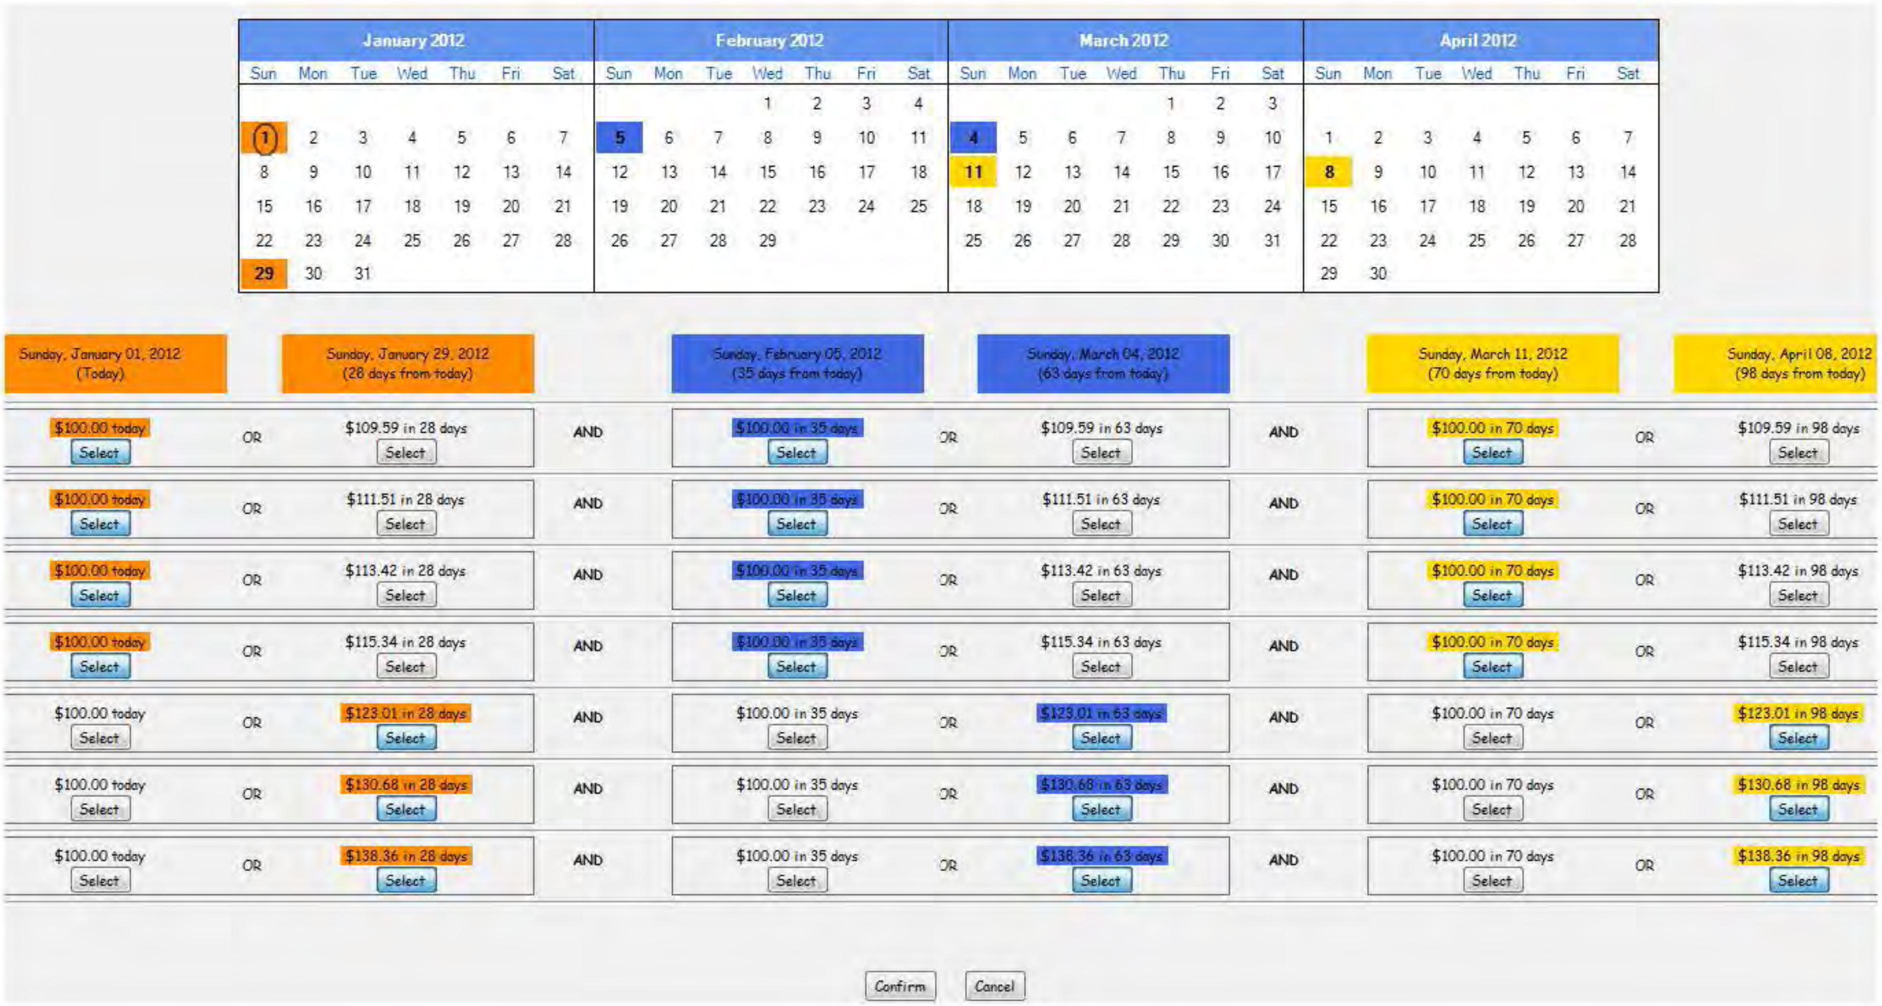
January 2012
February 2012
March 2012
April 2012
Sun Mon Tue Wed Thu Fri Sat
Sunday, January 01, 2012 (Today)
Sunday, January 29, 2012 (28 days from today)
Sunday, February 05, 2012 (35 days from today)
Sunday, March 04, 2012 (63 days from today)
Sunday, March 11, 2012 (70 days from today)
Sunday, April 08, 2012 (98 days from today)
$100.00 today Select OR $109.59 in 28 days Select AND $100.00 in 35 days Select OR $109.59 in 63 days Select AND $100.00 in 70 days Select OR $109.59 in 98 days Select
$100.00 today Select OR $111.51 in 28 days Select AND $100.00 in 35 days Select OR $111.51 in 63 days Select AND $100.00 in 70 days Select OR $111.51 in 98 days Select
$100.00 today Select OR $113.42 in 28 days Select AND $100.00 in 35 days Select OR $113.42 in 63 days Select AND $100.00 in 70 days Select OR $113.42 in 98 days Select
$100.00 today Select OR $115.34 in 28 days Select AND $100.00 in 35 days Select OR $115.34 in 63 days Select AND $100.00 in 70 days Select OR $115.34 in 98 days Select
$100.00 today Select OR $123.01 in 28 days Select AND $100.00 in 35 days Select OR $123.01 in 63 days Select AND $100.00 in 70 days Select OR $123.01 in 98 days Select
$100.00 today Select OR $130.68 in 28 days Select AND $100.00 in 35 days Select OR $130.68 in 63 days Select AND $100.00 in 70 days Select OR $130.68 in 98 days Select
$100.00 today Select OR $138.36 in 28 days Select AND $100.00 in 35 days Select OR $138.36 in 63 days Select AND $100.00 in 70 days Select OR $138.36 in 98 days Select
Confirm
Cancel

This screen shows seven triples of tied-together decisions. Each decision triple is presented on a different row. All decision triples have the same format. For the purpose of explaining this task, assume for the moment that today is January 1, 2012. In the calendar, today's date is enclosed in a black circle. Let's look at the first decision triple in the example above (the one on the first decision row). In the first decision, the sooner choice pays $100 today, January 1, 2012 in the example, and the later choice pays $109.59 in twenty-eight days from today. The second decision in this row offers the choice between $100 in thirty-five days or $109.59 in sixty-three days from today. The third decision in this row offers the choice between $100 in seventy days or $109.59 in ninety-eight days from today.

If you choose the sooner option for one decision in a triple, you are also choosing the sooner option for the other two decisions in that triple. In other words, it is not possible to choose the sooner option in one decision and the later option in the other decisions of the same triple. If you select the sooner options in one triple, you are free to select the later options in any other triple. You choose by clicking the button under the alternative you prefer.

We will present you with ten of these decision screens, with each screen displaying seven choices for you to make. You must make all seven choices on each decision screen before moving to the next decision screen. While on a single decision screen, the only difference between decision triples is that the dollar amounts of the future payment will change. However, triples on different decision screens will have different dollar amounts and future payment dates. So, you should make sure to pay attention to both the changing dollar amounts and changing dates as you make your decisions.

After you have worked through all of the decisions, we will select one of your ten decision screens by rolling a 10-sided die. Then we will roll a 10-sided die again, until a number between 1 and 7 comes up, to pick one decision row on that screen. This means you will be paid for all three choices in that triple, which you made at the same time. When you make your choices you will not know which decision triple will be selected for payment. You should therefore treat each decision as if it might actually count for payment.

Once the decision screen and row are selected, we will look at the specific choices you made and the resulting payouts. If the individual payouts in the selected decision row are smaller (less than $100), you will actually be paid all three amounts. However, if the individual payouts in the selected decision row are larger ($100 or greater), you will roll a 10-sided die to determine whether or not you are actually paid these amounts. If you roll a 1, you will actually be paid the amounts, and on the dates that you chose to receive them. If you roll a number other than 1, you will earn nothing in this task. This roll will be at the end of the experiment, and in private, when you are being paid.

For instance, suppose the decision screen in the above example was selected and you preferred the sooner date in the first four rows and the later date in the last three rows, as shown above. You would then roll a 10-sided die, until a number between 1 and 7 comes up, to select the row. Suppose the outcome was 5 and the fifth row is selected. In this row you preferred a payment of $123.01 in twenty-eight days in the first decision and, consequently, $123.01 in sixty-three days in the second decision, and $123.01 in ninety-eight days in the third decision. Since this payment is $100 or greater, you would roll the 10-sided die again, at the end of the experiment. If the outcome of your die-roll is 1, you will be paid $123.01 in twenty-eight days, $123.01 in sixty-three days, and another $123.01. However, if the outcome is 3, or anything other than 1, you will get nothing.

You will receive the money on the date stated in your preferred option. We will pay you using *PayPal*, which is an online payment service. We will explain more about *PayPal* in a few minutes. If you receive some money to be paid in the future you will also receive a written confirmation from Professor Harrison which guarantees that the money is to be paid to you on that date.

The money you receive from these choices is in addition to the show-up fee of $5, which is paid out at the end of the experiment as cash.