Intertemporal Consumption with Risk: A Revealed Preference Analysis

Lanier, Joshua and Miao, Bin and Quah, John and Zhong, Songfa

Oxford University, Shanghai University of Finance and Economics, Johns Hopkins University, National University of Singapore

1 February 2018

Intertemporal Consumption with Risk: A Revealed Preference Analysis

Joshua Lanier, Bin Miao, John K.-H. Quah, Songfa Zhong

Abstract

We run an experiment designed to elicit preferences over state contingent, timed payouts. We analyze the data using a new revealed preference method (building on Nishimura, Ok, and Quah (2017)) that can test for consistency with utility functions that increase with a given preorder. Using this approach, we find strong evidence of correlation averse behavior, a property ruled out by discounted expected utility. There is also evidence in favor of stochastic impatience.

Keywords: correlation aversion, temporal risk aversion, discounted expected utility, stochastic impatience, nonparametric methods

1 Introduction

I returned, and saw under the sun, that the race is not to the swift, nor the battle to the strong . . . but time and chance happeneth to them all.
(Ecclesiastes 9:11)[1]

As the writer of Ecclesiastes acknowledges, time and risk are abiding features of human existence. And yet easy answers are in short supply when modeling decision making where time and risk are involved. The simplest approach, discounted expected utility (DEU) is widely-acknowledged to be too restrictive for many purposes, yet exactly how DEU should be relaxed is not fully resolved.

The contribution of this paper is twofold. First, we contribute to the understanding of behavior under time and risk by conducting and analyzing the results of an experiment where participants make investment decisions over risky assets that pay out on different dates. Our analysis convincingly rejects a property of preferences called correlation

[1] This passage from the King James version of the Bible was given a translation by George Orwell into what he called "modern English of the worst sort": "Objective considerations of contemporary phenomena compel the conclusion that success or failure in competitive activities exhibits no tendency to be commensurate with innate capacity, but that a considerable element of the unpredictable must invariably be taken into account."

neutrality, which is a major consequence of the canonical DEU model; instead, we find evidence in favor of correlation aversion and also of stochastic impatience.

Our analysis is carried out with a new toolkit of techniques from the revealed preference literature, by operationalizing the results in Nishimura, Ok, and Quah (2017). These techniques, which are potentially applicable to other choice environments, allow us to test various structural properties of a subject's preference, without imposing functional form assumptions. Thus this paper also makes a second, methodological, contribution to the use of revealed preference techniques in empirical research.

1.1 The Experimental Setup

We task 103 subjects with making 41 different decisions. Each decision involves spending 100 tokens to purchase quantities of four different assets. There are two equally likely states (call them state 1 and state 2) and each of the four assets delivers a payout in only one of the two states and at one of two dates. Thus, each asset in our experiment is characterized by a pair (s, t) where $s \in \{1, 2\}$ represents the state of the world in which the asset pays out money and $t \in \{1, 2\}$ represents the time of payment (1 being earlier and 2 being later). We vary the prices of the assets, giving rise to 41 decision problems.

How participants choose in these decisions convey information on their preferences. For example, a participant may allocate their entire budget to the asset with the highest payout, which would signal extreme risk neutrality. At another extreme, a participant may purchase the same amount of each asset, thus ensuring the same payout in each state and at each time period. Importantly, the participant may also signal their level of aversion to risks correlated across time. For example, a participant who purchases many units of an asset which pays out in state 1 period 1 shall perhaps prefer to purchase the asset which pays out in state 2 period 2 when compared to the asset which pays out in state 1 period 2. In this way the participant may signal their desire for or aversion to payouts correlated over time.

In many experiments, subjects are asked which of two options they prefer. By changing the options offered the researcher gets a sense of the preferences of the participant. In contrast, subjects choose from a linear budget set in our experiment; an early example of this type of experiment appears in Choi, Fisman, Gale, and Kariv (2007) where risk preferences are explored.[2] Besides its resemblance to 'real life' decision-making, budgetary decisions have the advantage that they convey information efficiently, in the sense that the choice of one bundle from a budget means that it is preferred to an infinite set

[2]Linear budgets are also used to investigate risk preferences experimentally in Choi, Kariv, Wieland, and Silverman (2014), Ahn, Choi, Gale, and Kariv (2014), Halevy, Persitz, and Zrill (2018), and time preferences in Andreoni and Sprenger (2012b).

of alternatives. The richer information makes it possible for the researcher to estimate each participant's preference separately (in other words, allowing for full heterogeneity) and to test a large number of hypotheses about a participant's preference (beginning with whether or not the participant is maximizing some utility function, and going on to various restrictions on its structure). Since the experimental setting is common to all hypotheses tested it also facilitates comparisons among them.

1.2 Revealed Preference Analysis

We analyze the data with a revealed preference approach. The canonical result of Afriat (1967) (see also Afriat (1972)) provides a simple non-parametric test for the hypothesis that a dataset of budgetary choices is generated by a utility-maximizing agent. The test is non-parametric in the sense that there are no functional form restrictions on the utility function and it is merely required to be increasing (more is better). The outcome of the test is a number $e \in [0, 1]$, referred to as the *critical cost efficiency index* (or just the efficiency level for short), which represents the goodness-of-fit of the utility maximization hypothesis. A value of 1 corresponds to a perfect fit, wherein the data could be exactly generated by utility maximization. If e is less than 1 then we may conclude that if the dataset was generated by a utility maximizer, then this agent must, at some observations, waste a fraction $1 - e$ of their budget set, in the sense that the same utility can be achieved by spending a fraction e of the actual amount spent. An agent with a high e is close to being consistent with utility maximization.

In this paper we are interested in testing if the choices from a participant in the experiment could have come about as the result of maximizing a utility function satisfying some additional property $\mathbf{P}$. To implement tests of this type we exploit an extension of Afriat's Theorem found in Nishimura, Ok, and Quah (2017) (henceforth NOQ). Not every property $\mathbf{P}$ is covered by the NOQ approach, but – crucially – it *does* cover some of the core properties we care about, such as correlation aversion, impatience, and stochastic impatience. For example, we can calculate the efficiency level at which a participant is consistent with (basic) utility maximization and also the efficiency level at which the participant is consistent with the maximization of a utility function that satisfies impatience. (By definition the latter will be weakly lower than the former.)

1.3 The Results

The DEU model assumes that a temporal sequence of random variables $X = (X^0, \ldots, X^T)$ representing monetary payouts is valued by a utility function of the form

$$U(X) = \mathbf{E}\left[\sum_{t=1}^{T} \delta^t u(X^t)\right] \tag{1}$$

for some increasing function $u : \mathbb{R}_+ \to \mathbb{R}$. The first application of our methodology imposes two (we believe) uncontroversial properties on to the utility function; termed lottery equivalence and impatience. Both properties are implied by the DEU model.

Lottery equivalence says that the agent only cares about the distribution of payouts and not the names of the states in which payouts occur. Concretely it says that assets which pay out in state 1 are just as good as assets which pay out in state 2. **Impatience** is the familiar preference for sooner payouts compared to later. This first exercise may be thought of as a basic sense check of our data and methodology. While it should be unsurprising to confirm these properties we view these first exercises as an important way of ensuring that participants were not pushed beyond their cognitive capacities by the experiment. If we could find little evidence in support of these two uncontroversial properties then it would make us doubt if we could draw meaningful conclusions concerning less established properties. Thankfully, we find strong evidence that the participants in the experiment are utility maximizers satisfying these two properties. In particular we find that almost 80% of the participants act as though they are maximizing utility functions satisfying lottery equivalence and impatience at the 0.95 efficiency level. That is, they never waste more than 5% of their budgets.

We next test whether participants satisfy the property of **correlation neutrality**, another property implied by DEU. Correlation neutrality means that agents are indifferent to changes in the correlation between payouts received at different time periods. To illustrate let us compare the following gambles where, as in our experiment, we assume there are two equiprobable states. Gamble 1 pays out \$1 in period 1 and \$1 in period 2 should state state 1 occur and pays out nothing should state 2 occur. Gamble 2 pays out \$1 in period 1 and \$0 in period 2 should state 1 occur while it pays out \$0 in period 1 and \$1 in period 2 should state 2 occur. With Gamble 1 a nice payout sooner is accompanied by a nice payout later. On the other hand with Gamble 2 a good payout sooner is accompanied by a bad payout later. Under correlation neutrality a person is indifferent between Gamble 1 and Gamble 2. **Correlation aversion** is the preference for payouts across time to be negatively correlated and so a correlation averse agent prefers Gamble 2 to Gamble 1. It is straightforward to see that the DEU formula (1) satisfies correlation neutrality.[3]

We find strong evidence that **participants do not display correlation neutrality and instead display correlation aversion**. We find that a mere 22.3% of our participants have efficiency levels in excess of 0.95 when assuming their preferences exhibit correlation neutrality. What this means is that if one wishes to insist that our participants truly exhibit correlation neutrality then one must be willing to accept that more than 75% of our participants routinely waste more than 5% of their budget. In contrast, if we impose lottery equivalence, impatience, and correlation aversion then we find that more than 75% of our participants have an efficiency level in excess of 0.95.

A second application of our revealed preference methods finds evidence in favor of **stochastic impatience**. We believe the present study is the first to test for this property. We again illustrate this property by comparing two gambles. Gamble 3 pays out either 2 dollars in period 1 or 1 dollar in period 2 each with 50% probability. Gamble 4 pays out either 1 dollar in period 1 or 2 dollars in period 2 each with 50% probability. The stochastically impatient agent prefers Gamble 3 to Gamble 4. That is, the agent prefers to have a chance of receiving the good 2 dollar payment earlier to the scenario where the good payment is received in the later period. This property is studied in DeJarnette, Dillenberger, Gottlieb, and Ortoleva (2019) and Dillenberger, Gottlieb, and Ortoleva (2018). We find every subject who has an efficiency level of 0.95 for the combined properties of impatience and lottery equivalence (near 80% of all subjects) also has an efficiency level above 0.95 when requiring stochastic impatience as well. Thus, stochastic impatience performs well though some caution may be in order for reasons given in Section 5.4.

1.4 Contribution to revealed preference techniques

Our revealed preference methodology relies on applying Theorem 2 in NOQ. This result allows one to test if choices from a consumer are consistent with the maximization of a utility function U satisfying a property $\mathbf{P}$, where $\mathbf{P}$ can be represented by some ordering $\trianglerighteq$.[4] The properties we test in this paper can be so represented and thus the NOQ result applies in principle.

To be specific, suppose B^t is the linear budget set faced by a participant in period t and suppose $c^t \in B^t$ is the choice made by the participant in period t. To utilize the NOQ Theorem one must, for every observation number s and t, ascertain whether there

[3] The term 'correlation aversion' comes from Epstein and Tanny (1980). This property is referred to as multivariate risk aversion in Richard (1975) and temporal risk aversion in Karady (1982).

[4] Formally, $\trianglerighteq$ is required to be a preorder. $\trianglerighteq$ represents some property $\mathbf{P}$ if U satisfies $\mathbf{P}$ if and only if the following holds: $U(c) \geq U(\bar{c})$ whenever $c \trianglerighteq \bar{c}$.

exists some $\bar{c} \in B^t$ so that $\bar{c} \trianglerighteq c^s$. As B^t has infinitely many elements, checking this need not be computationally feasible. In this article we identify conditions on $\trianglerighteq$ which ensure that there is only a *finite* number $\bar{c}$ for which the check needs to be carried out. In this way, we are able to operationalize NOQ's Theorem 2 and apply it to a large and relevant class of properties **P**.

1.5 Remainder of the Paper

The remainder of the paper is organized as follows. Section 2 provides a review of related literature; readers anxious to get to the results can skip this in the first reading. Section 3 presents our experimental design and a preliminary description of our data. Section 4 explains our revealed preference tests. Section 5 presents the results. Section 6 provides additional tests, including some where a parametric or semiparametric form are assumed. This is followed by an Appendix containing proofs and other material.

2 Related Literature

The DEU model is the canonical approach to modeling consumer behavior when household decisions are made in a risky environment over multiple periods. Andersen, Harrison, Lau, and Rutström (2008) show how features of the DEU model can be leveraged to recover the discount factor δ using an experiment designed to separately elicit preferences for risk (allowing for identification of u) and preferences for non-risky intertemporal payouts (allowing identification of δ).[5]

The DEU model implies that risky prospects whose payouts are all delivered in the same period are evaluated using the Expected Utility (EU) model and that non-risky temporal streams of payouts are evaluated using the Discounted Utility (DU) model. Thus, the DEU model inherits any undesirable or unrealistic properties of both the EU and the DU models. For example, the EU model cannot account for the common consequence or common ratio effects and thus neither can DEU (see Starmer (2000)).[6] The DU model cannot account for the common difference effect and thus neither can DEU (see Frederick, Loewenstein, and O'Donoghue (2002)).[7]

[5]The case of hyperbolic discounting is also considered. A similar exercise is pursued in Ferecatu and Öncüler (2016) using Bayesian methods.

[6]Let $(a_1, a_2; \pi, 1-\pi)$ refer to lottery which pays out a_1 with probability π and a_2 with probability $1-\pi$. A common consequence effect refers to a reversal in preference $(a_1, a_2; \pi, 1-\pi) \succ (\ell, a_2; , \pi, 1-\pi)$ when a_2 is changed. Here ℓ refers to some non-degenerate lottery and a_2 is the "common consequence". A common ratio violation refers to a reversal in preference $(a_1, 0; \pi, 1-\pi) \succ (a_2, 0; \pi', 1-\pi')$ when π and π' are scaled by some common factor $\alpha \in (0, 1)$.

[7]A common difference effect is present when a preference reversal is observed when adding a common

There are also experiments which show that *combining* risk and time (for example, with lotteries which pay out over multiple periods) can produce interesting and unexpected behavior. Keren and Roelofsma (1995) show that adding risk to timed payouts can negate the common difference effect.[8] Baucells and Heukamp (2010) show that adding a delay to lotteries can trigger a version of the common ratio effect. Halevy (2008), Epper and Fehr-Duda (2018), and Chakraborty, Halevy, and Saito (2020) explain experimental findings on certain departures from EU, DU, and DEU by applying the rank dependent utility model in Quiggin (1982).[9,10] Andreoni and Sprenger (2012a) present experimental results demonstrating that subjects display a version of the common ratio effect that contradicts the DEU model.[11]

The DEU model requires that the function u used to evaluate risk is the same as that used to evaluate time. This implies that the coefficient of risk aversion is the inverse of the coefficient of intertemporal substitution. Coble and Lusk (2010), Andreoni and Sprenger (2012a), and Abdellaoui, Bleichrodt, L'Haridon, and Paraschiv (2013) present experimental evidence showing that these two functions are not the same.[12]

The DEU model imposes correlation neutrality. Kihlstrom and Mirman (1974), Richard (1975), Epstein and Tanny (1980), Bommier (2007), and Van den Heuvel (2008) present theoretical approaches to modeling attitudes towards intertemporal risk that allows for correlation non-neutrality. The empirical evidence on the validity of correlation neutrality is less developed. A recent contribution is Andersen, Harrison, Lau, and Rutström (2018), who implement an experiment (involving binary choices) on a representative sample of the Danish population. They use this data to estimate the utility functions of representative agents for different segments of the population; these functions are assumed to take a parametric form within the Kihlstrom-Mirman class (see Kihlstrom and Mirman (1974)). They find that this agent is correlation averse and their approach also allows them to measure the degree of correlation aversion. Our findings on correlation aversion are broadly consistent with theirs, even though our experiment and analysis are different. The budgetary data we collect allows us to estimate nonparamet-

delay to the problem of choosing between receiving prize A in period t and prize B in period t'.

[8]Weber and Chapman (2005) largely confirm these results but with some nuance.

[9]Other experiments documenting interesting interactions between risk and time preferences can be found in Anderhub, Güth, Gneezy, and Sonsino (2001), Onay and Öncüler (2007), Epper, Fehr-Duda, and Bruhin (2011), and Abdellaoui, Kemel, Panin, and Vieider (2019).

[10]See also Baucells and Heukamp (2012) for a model which can explain some of the departures from EU and DU.

[11]This paper has given rise to interesting follow up discussion. See Cheung (2015), Miao and Zhong (2015), Epper and Fehr-Duda (2015), Andreoni, Feldman, and Sprenger (2017).

[12]Various theoretical models have been developed that allow for a relaxing of the requirement that risk and time are evaluated using the same function u. Examples include Selden (1978), Kreps and Porteus (1978), Epstein and Zin (1989), Chew and Epstein (1990), and Kubler, Selden, and Wei (2019).

rically (and separately) the utility function of each participant in the experiment and to ascertain if correlation aversion holds for a given participant.

It is shown in DeJarnette et al. (2019) that DEU implies a property they call risk seeking over time lotteries (RSTL). RSTL says that a guaranteed payment of receiving $x > 0$ dollars with 50 percent chance in period 1 and 50 percent chance in period 3 is preferred to a guaranteed payment of x dollars in period 2. An experiment presented in DeJarnette et al. (2019) provides evidence that participants often violate RSTL, thus demonstrating another problem with the DEU model.

3 Experiment Design and Summary Statistics

We provide subjects with a budget of 100 experimental tokens which they can allocate over four state and time contingent commodities. A typical consumption bundle is denoted by

$$c = (c_{11}, c_{12}, c_{21}, c_{22})$$

where $c_{st} \geq 0$ refers to consumption in state s at time t. The two states, determined by a coin toss, are of equal probability; and the two time points are one week later and 9 weeks later. The price vector faced by the subject is denoted by

$$p = (p_{11}, p_{12}, p_{21}, p_{22})$$

where $p_{st} > 0$ refers to the price of a unit of consumption in state s at time t. In the experiment, p is generated by having the price for each commodity p_{st} randomly selected from the set $\{0.5, 0.8, 1, 1.25, 2\}$, with at least one price being equal to 1. This gives rise to 96 distinct price vectors. We randomly select 41 price vectors for each subject with one of them fixed to be the benchmark vector $(1, 1, 1, 1)$.

At the end of the experiment, each subject is paid according to one randomly selected decision task by tossing dice according to the Random Incentive Mechanism (RIM).[13] Subjects are informed to treat each decision as if it is the sole decision determining their payments. To control for preferences over the timing of uncertainty resolution, all uncertainty is resolved at the end of the experiment. Once the decision task is selected and the state is realized, each subject is paid with an exchange rate of SGD 0.2 (about USD 0.15) per experimental token. To increase the credibility of payment, subjects are paid with post-dated checks that will not be honored by the local bank when presented prior to the date indicated. To further control for the potential difference in transaction cost

[13]For the validity of RIM, readers can refer to Wakker (2007) for a detailed discussion.

at different time points (Andreoni and Sprenger (2012b)), subjects receive a minimum participation fee of SGD 12 (with SGD 6 for each payday). Experimental earnings are added to these minimum payments.

While our experiment uses monetary rewards, utility is derived from actual consumption in most theoretical models.[14] Halevy (2014) provides a detailed defense of the use of monetary rewards in the study of time preference. Among other things, he highlights the considerable evidence in favor a high correlation between modest income gains (such as those won in an experiment) and contemporaneous consumption; this suggests that concerns about the use of monetary rewards, which are typically based on some version of the permanent income hypothesis, may be overstated. It is also worth observing that our experiment involves both time and risk, and those experiments which involve only risk or uncertainty are often based on monetary rewards. Thus our use of monetary rewards makes it easier for us to compare our conclusions on subjects' risk attitudes with those conclusions drawn from experiments that involve only risk.

A total of 103 undergraduate students were recruited as participants through an advertisement posted in the Integrated Virtual Learning Environment at the National University of Singapore. The experiment was conducted at the laboratory of the Center for Behavioral Economics at the National University of Singapore. Conducted by two of the authors and a research assistant, the experiment consists of four sessions with 20 to 30 subjects in each session. After the subjects arrived at the experiment venue, they were given the consent form approved by the Institutional Review Board of the National University of Singapore. Following that, general instructions were read aloud to the subjects, and several examples were demonstrated to them before they started making their decisions. The experimental instructions follow closely those in Andreoni and Sprenger (2012a) (See the Online Appendix for the Experimental Instructions). Most of our subjects completed the tasks within 30 minutes. At the end of the experiment, they approach the experimenters one by one, toss the dice and receive payments in post-dated checks based on their choice. On average, the subjects are paid SGD 22.

[14] Augenblick, Niederle, and Sprenger (2015) examine time preference over effort and show that the incentive for smoothing is higher for effort than for monetary rewards. Nonetheless, Reuben, Sapienza, and Zingales (2010) elicit discount factors for both monetary rewards and primary rewards of chocolate and find a positive and statistically significant relation between the two. The observed correlation is further corroborated by neuroimaging studies which show that the same limbic areas of the brain are activated irrespective of monetary or primary rewards (McClure, Laibson, Lowenstein, and Cohen (2004); McClure, Ericson, Laibson, Loewenstein, and Cohen (2007)).

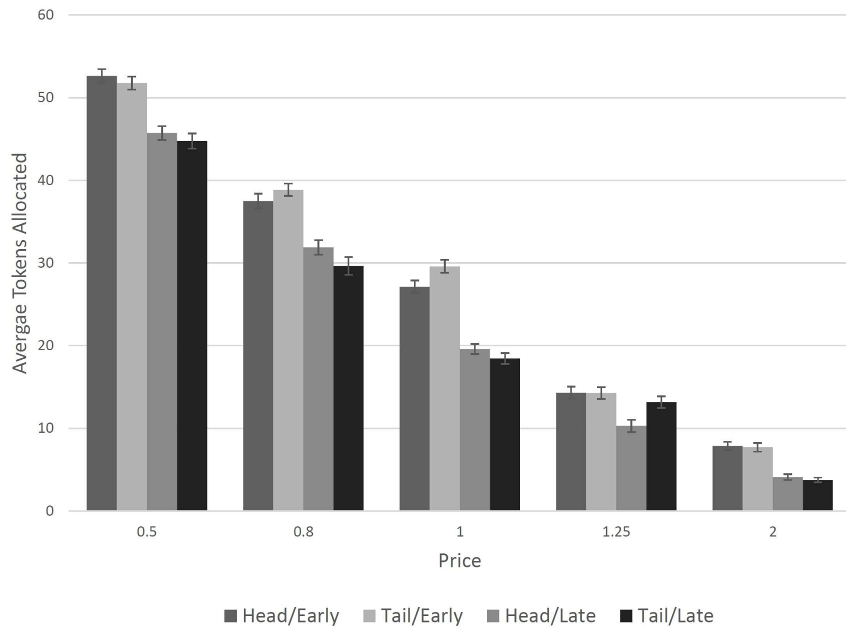

Figure 1: Average tokens allocated to each asset

3.1 Aggregate Behavior

Figure 1 plots the average allocation of tokens for each of the four commodities under each price.[15] Two patterns arise. First, the average allocation is lower when the price is higher, suggesting that the law of demand is satisfied in aggregate. Second, at any given price, the allocation to the early time point is larger than that to the late time point, which is indicative evidence that subjects are on average impatient.

We further conduct regression analyses with the tokens allocated to each commodity as dependent variables and the prices for all the commodities as independent variables. We apply a Tobit regression model with censoring at both 0 and 100, given the concern of corner choices. From the results reported in Table 1 below, we observe that the tokens allocated to each commodity is negatively affected by its own price, and positive affected by the price of the other three commodities. Moreover, the cross-price effect is stronger if the two commodities are within same state, compared to when they are within the same time period. For example the price effect of state 1 time 2 price changes on the state 1 time 1 consumption is 24.757. This is much larger than the price effect of state 2 time 1 price changes on the state 1 time 1 consumption which is 3.660. This suggests that the motive for diversification across states is stronger than the motive for

[15]This is computed as the average allocation to an asset at a given price for that asset, across all observations and all subjects, without regard to the prices of other assets.

smoothing across time points.

Table 1: Regression Analysis on Price Effect.

	$c_{1,1}$	$c_{2,1}$	$c_{2,1}$	$c_{2,2}$
$p_{1,1}$	**-38.831*** **	3.902***	31.237***	2.367**
	(3.872)	(1.425)	(2.584)	(1.157)
$p_{2,1}$	3.660***	**-38.384*** **	4.640***	31.198***
	(1.205)	(3.497)	(1.458)	(2.715)
$p_{1,2}$	24.757***	3.517***	**-40.866*** **	1.928
	(2.138)	(1.300)	(4.447)	(1.302)
$p_{2,2}$	1.924**	23.524***	3.750***	**-39.549*** **
	(0.862)	(1.821)	(1.305)	(4.317)
constant	30.120***	28.292***	9.147**	12.055***
	(2.846)	(3.237)	(3.737)	(3.401)
Observations	4,223	4,223	4,223	4,223
Pseudo R-squared	0.0659	0.0689	0.0701	0.0745

We find a substantial number of corner choices made. By a corner choice we mean a round of the experiment where the subject allocates 0 tokens to one of the assets. We find that 49 percent of participants make corner choices in all rounds of the experiment and 91 percent make at least one corner choice. These numbers seem consistent with convex time budget experiment of Andreoni and Sprenger (2012b) where 37 percent of subjects make only corner solutions.[16]

4 Utility Functions and Revealed Preference Tests

The goal of this section is to explain the revealed preference techniques used in our analysis. We shall only be interested in utility functions which are continuous and increasing and so we embed this into the definition of utility.

Definition 1. A *utility function* is a strictly increasing and continuous function $U : \mathbb{R}^4_+ \to \mathbb{R}$.

4.1 Quantifying Imperfect Utility Maximization

An *e-utility maximizer* is an agent who possesses some utility function U and some number $e \in [0, 1]$ so that when choosing from budget set $B(p, w) = \{c \in \mathbb{R}^4_+ : p \cdot c \leq w\}$, selects a bundle $\bar{c} \in B(p, w)$ that satisfy

$$U(\bar{c}) \geq U(c), \quad \text{for all } c \in B(p, ew) \tag{2}$$

[16]It seems reasonable that simply having more commodities would make corner solutions more likely. The experiment of Andreoni and Sprenger (2012b) has only two goods (there is no risk) and so the presence of fewer all-corner solutions may be due to the fact that they have fewer goods.

Such an e is said to be the agent's *efficiency level.* Of course, an agent who maximizes U with $e = 1$ is a utility-maximizing agent in the familiar sense.

Given a data set of demand observations, we could ask if it is consistent with utility-maximization at a given efficiency level. This is captured by the following definition.

Definition 2. Suppose $c^1, \ldots, c^N$ are the bundles chosen by a subject in rounds 1 through N of the experiment and let $p^1, \ldots, p^N$ be the corresponding price vectors. We refer to the collection $\mathcal{O} = \{(c^1, p^1), \ldots, (c^N, p^N)\}$ as a *dataset.* The dataset $\mathcal{O}$ is *rationalized* by utility function U at efficiency level e if the choices made could have been made by an e-utility maximizer with utility function U. That is, for all n

$$U(c^n) \geq U(c), \qquad \text{for all } c \in B(p^n, ep^n \cdot c^n) \tag{3}$$

A dataset $\mathcal{O}$ that admits such a utility function U is said to be *e-rationalizable.*

Obviously, if a dataset $\mathcal{O}$ is e-rationalizable then it is also e'-rationalizable for any $e' < e$. The *critical cost efficiency index* (see Afriat (1972)) is defined as

$$\bar{e} \;=\; \sup \Big\{ e \in [0, 1] : \mathcal{O} \text{ is } e\text{-rationalizable} \Big\}. \tag{4}$$

4.2 Utility functions satisfying further properties

Let $\mathbf{P}$ denote a property which a utility function may or may not possess[17] and let $e \in [0, 1]$. Consider the choices of one subject from our experiment. We would like to develop a way to answer the following question: does there exist a utility function U satisfying $\mathbf{P}$ that rationalizes the choices of the subject at efficiency level e? We focus on properties $\mathbf{P}$ which can be represented by a preorder.[18]

Definition 3. Let $\trianglerighteq$ be a preorder and let U be a utility function. Utility function U *agrees* with $\trianglerighteq$ if

$$c \trianglerighteq c' \qquad \Longrightarrow \qquad U(c) \geq U(c'), \qquad \text{for all } c,\ c' \in \mathbb{R}^4_+ \tag{5}$$

While it can be difficult to determine if a subject's choices can be rationalized by a utility function that agrees with some arbitrary preorder $\trianglerighteq$, we shall see that when the preorder satisfies the following property the task is greatly simplified.

[17] Formally, one may take $\mathbf{P}$ to be a set of utility functions. Then, when we say U possesses property $\mathbf{P}$ we just mean $U \in \mathbf{P}$.

[18] A preorder is a transitive and reflexive binary relation.

Definition 4. A preorder $\trianglerighteq$ is a *composable preorder* if for all $c, c' \in \mathbb{R}^4_+$

$$\text{there exists } \bar{c} \text{ so that } c \trianglerighteq \bar{c} \geq c' \qquad \Longleftrightarrow \qquad \text{there exists } \tilde{c} \text{ so that } c \geq \tilde{c} \trianglerighteq c'$$

where $\geq$ is the product order on $\mathbb{R}^4_+$.

Note that, by definition a utility function is an increasing function and thus it must agree with the product order $\geq$; what composability guarantees is that the added property $\trianglerighteq$ interacts nicely with $\geq$. We impose two other, relatively innocuous conditions (in addition to the composable condition), on all preorders considered in this paper and refer to a preorder which satisfies all three conditions as *well-behaved*. A full definition of a well-behaved preorder can be found in the Appendix (Section A.1).

We now present two key preference properties studied in this paper; each property is captured by agreement with a particular well-behaved preorder. It is convenient to denote the consumption bundle of an agent by $c = (c_1, c_2)$, where $c_1 = (c_{11}, c_{12})$ denotes state 1 consumption (at dates 1 and 2) and $c_2 = (c_{21}, c_{22})$ denotes state 2 consumption.

Lottery Equivalence This property says that swapping the consumption streams in states 1 and 2 makes no difference to utility. It makes sense if the realized state has no intrinsic significance for the agent's preference. Thus the bundles (a, b) and (b, a) would both simply correspond to a lottery that pays out streams a and b with equal probability and thus must give the same utility.

Definition 5. A utility function $U : \mathbb{R}^4_+ \to \mathbb{R}$ satisfies *lottery equivalence* if

$$U(a, b) = U(b, a), \qquad \text{for all } a, b \in \mathbb{R}^2_+.$$

The lottery equivalence property can be re-stated as an agreement with a preorder $\trianglerighteq$. The lottery equivalence preorder $\trianglerighteq_{LE}$ is defined as follows:

$$(a, b) \trianglerighteq_{LE} (a, b) \qquad \text{and} \qquad (a, b) \trianglerighteq_{LE} (b, a)$$

for all $a, b \in \mathbb{R}^2_+$. It is straightforward to check that $\trianglerighteq_{LE}$ is a preorder and that it is composable (and in fact it is also well-behaved). The asymmetric part of $\trianglerighteq_{LE}$ is empty: there are no elements c and c' such that $c \triangleright_{LE} c'$. The lottery equivalence property for a utility function U can be equivalently stated as follows: $U(c) \geq U(c')$ whenever $c \trianglerighteq_{LE} c'$.

Impatience. In models of time preference, it is common to postulate that agents are impatient, which means that an agent has a preference for receiving sooner payments over later payments. When confined to a consumption stream with no risk (in our context, this would be a bundle of the form $((a,b),(a,b))$), the property says that $U((a,b),(a,b)) \geq U((b,a),(b,a))$ if $a \geq b$. A stronger version of this property is to require impatience to hold state-by-state, which is captured by the following definition.

Definition 6. The utility function $U : \mathbb{R}^4_+ \to \mathbb{R}$ satisfies *impatience* if

$$U((a,b),(a',b')) \geq U((\tilde{a},\tilde{b}),(\tilde{a}',\tilde{b}'))$$

whenever any of the following hold:
(i) $a \geq b$, $\tilde{a} = b$, $\tilde{b} = a$, and $(a',b') = (\tilde{a}',\tilde{b}')$;
(ii) $a' \geq b'$, $\tilde{a}' = b'$, $\tilde{b}' = a'$, and $(a,b) = (\tilde{a},\tilde{b})$; and
(iii) $a \geq b$, $a' \geq b'$, $\tilde{a} = b$, $\tilde{b} = a$, and $\tilde{a}' = b'$, $\tilde{b}' = a'$.[19]

We define $\trianglerighteq_I$ as follows: $((a,b),(a',b')) \trianglerighteq_I ((\tilde{a},\tilde{b}),(\tilde{a}',\tilde{b}'))$ if either $((a,b),(a',b')) = ((\tilde{a},\tilde{b}),(\tilde{a}',\tilde{b}'))$ or (i), (ii) or (iii) in Definition 6 holds. It is straightforward to check that $\trianglerighteq_I$ is composable (indeed, well-behaved) preorder with a nonempty asymmetric part,[20] and a utility function U satisfies impatience if and only if $U(c) \geq U(c')$ whenever $c \trianglerighteq_I c'$.

Lottery Equivalence and Impatience. Both lottery equivalence and impatience are natural properties in our choice environment and it would be useful to check if a dataset can be rationalized by a utility function with both properties. This corresponds to requiring that U agrees with the preorder $\trianglerighteq_{LEI}$, which is the transitive closure of $\trianglerighteq_{LE}$ and $\trianglerighteq_I$.[21] Our testing procedure is applicable so long as $\trianglerighteq_{LEI}$ is a well-behaved preorder. It turns out that that is indeed the case.

There are other properties studied in this paper which can be defined as agreement with a well-behaved preorder. Furthermore, these preorders can also be combined, via transitive closure, while remaining well-behaved (see Appendix (Sections A.2 and A.3) for the details). This means that we have the ability to test for e-rationalization by a utility function with any combination of such properties.

In our experiment the two states are equally likely. In this case, the utility function

[19](i) corresponds to the case where utility is increased if the state 1 consumption stream is altered from (b,a) to (a,b), which brings forward the larger consumption and postpones the smaller consumption, while keeping state 2 consumption fixed. Case (ii) corresponds to the case where the change takes place in state 2, while (iii) corresponds to the case where consumption in both states are altered.

[20]For example, $((2,1),c_2) \triangleright ((1,2),c_2)$ for any c_2.

[21]We write $c \trianglerighteq_+ c'$ if either $c \trianglerighteq_{LE} c'$ or $c \trianglerighteq_I c'$. By definition, $c \trianglerighteq_{LEI} c'$ if there exists some finite sequence $c^1,\ldots,c^K$ so that $c \trianglerighteq_+ c^1 \trianglerighteq_+ c^2 \trianglerighteq_+ \ldots \trianglerighteq_+ c^K \trianglerighteq_+ c'$.

$U : \mathbb{R}^4_+ \to \mathbb{R}$ has the *discounted expected utility* (DEU) form if

$$U(c_{11}, c_{12}, c_{21}, c_{22}) = \frac{1}{2}\left(u(c_{11}) + \delta u(c_{12})\right) + \frac{1}{2}\left(u(c_{21}) + \delta u(c_{22})\right) \tag{6}$$

for some strictly increasing function $u : \mathbb{R}_+ \to \mathbb{R}$ and $\delta \in (0, 1]$. It is clear that with this form, both lottery equivalence and impatience holds. However, this is by no means the only case where these properties are satisfied; we postpone the discussion of other functional forms with these properties to Section 4.4.

4.3 Revealed Preference and GARP

We refer to a dataset $\mathcal{O} = \{(c^1, p^1), \ldots, (c^N, p^N)\}$ as *e-rationalizable-*$\trianglerighteq$ if it can be rationalized by a utility function that agrees with $\trianglerighteq$ at efficiency level e. Clearly, if $\mathcal{O}$ is e-rationalizable-$\trianglerighteq$ for some e then $\mathcal{O}$ is e'-rationalizable-$\trianglerighteq$ for any $e' < e$. Our objective is to devise a way to calculate the (generalized) critical cost efficiency index

$$\bar{e}(\trianglerighteq) \;=\; \sup\Big\{ e \in [0, 1] : \mathcal{O} \text{ is } e\text{-rationalizable-} \trianglerighteq \Big\}. \tag{7}$$

Note that $\bar{e}$ as defined by (4) must satisfy $\bar{e} \geq \bar{e}(\trianglerighteq)$.

Let $I^{\trianglerighteq}(c)$ be the set of bundles ranked weakly better than c under the $\trianglerighteq$ preorder. That is

$$I^{\trianglerighteq}(c) = \{\bar{c} \in \mathbb{R}^4_+ : \bar{c} \trianglerighteq c\} \tag{8}$$

Given $\mathcal{O}$, let $\mathcal{C} = \{c^1, \ldots, c^N\}$. We define two binary relations on $\mathcal{C}$.

Definition 7. Let $\trianglerighteq$ be a preorder, $e \in [0, 1]$, and $\mathcal{O} = \{(c^1, p^1,), \ldots, (c^N, p^N)\}$ be a dataset. For c^n, $c^m \in \mathcal{C}$ write $c^n \trianglerighteq^*_e c^m$ if there exists $\bar{c} \in I^{\trianglerighteq}(c^m)$ so that

$$ep^n \cdot c^n \geq p^n \cdot \bar{c} \quad \text{ or } \quad c^n \geq \bar{c} \tag{9}$$

If $\geq$ can be replaced by its strict version in (9) then we write $c^n \triangleright^*_e c^m$.[22] We refer to $\trianglerighteq^*_e$ and $\triangleright^*_e$ as, respectively, the *revealed preferred* and *strictly revealed preferred* relations.

If a subject is truly an e-utility maximizer with a utility function that agrees with preorder $\trianglerighteq$, then $c^n \trianglerighteq^*_e c^m$ implies $U(c^n) \geq U(c^m)$ and $c^n \triangleright^*_e c^m$ implies $U(c^n) > U(c^m)$. Thus, the following property must hold for such a subject.

[22]Some authors use $\triangleright^*_e$ to represent the asymmetric part of the preorder $\trianglerighteq^*_e$. That is $c \triangleright^*_e c'$ if $c \trianglerighteq^*_e c'$ and not $c' \trianglerighteq^*_e c$. Note that our notation differs from this usage.

Definition 8. A dataset $\mathcal{O} = \{(c^1, p^1,), \ldots, (c^N, p^N)\}$ satisfies the *e-Generalized Axiom of Revealed Preference according to* $\trianglerighteq$ (e-GARP-$\trianglerighteq$)[23] if the revealed preference relations contain no strict cycles in the following sense: for any sequence $n_1, n_2, \ldots, n_K$ where

$$c^{n_1} \trianglerighteq_e^* c^{n_2} \trianglerighteq_e^* \ldots \trianglerighteq_e^* c^{n_K}$$

it is not the case that $c^{n_K} \triangleright_e^* c^{n_1}$.

The following proposition is the theoretical result underpinning the testing procedures in this paper. It relates the notion of e-GARP-$\trianglerighteq$ to that of rationalizing $\mathcal{O}$ with a utility function that agrees with $\trianglerighteq$.

Proposition 1. *Let* $\trianglerighteq$ *be a well-behaved preorder and let* $\mathcal{O} = \{(c^1, p^1), \ldots, (c^N, p^N)\}$ *be a dataset. Then* $\bar{e}(\trianglerighteq)$ *(as defined by (7)) satisfies*

$$\bar{e}(\trianglerighteq) = \sup \left\{ e \in [0, 1] : \mathcal{O} \textit{ satisfies } e\textit{-GARP-} \trianglerighteq \right\}$$

This proposition states that the critical cost efficiency index $\bar{e}(\trianglerighteq)$ equals the supremum of those efficiency levels e at which e-GARP-$\trianglerighteq$ holds. Determining the latter is straightforward once the revealed preference relations have been computed, since there are well known algorithms for checking the absence of strict cycles. Working out the revealed preference relations in turn require determining the set $I^{\trianglerighteq}(c)$ (defined in (8)). Crucially, for the preorders considered in this paper, this set is finite (see Appendix, Sections A.2 and A.3). For example, it is clear that

$$I^{\trianglerighteq_{LE}}(a, b) = \{(a, b), (b, a)\}.$$

Writing out $I^{\trianglerighteq_I}(a, b)$ or $I^{\trianglerighteq_{LEI}}(a, b)$ is a bit more longwinded but it is not difficult and, plainly, the set is finite. In the Appendix (Section A.4) we provide a systematic way of working out $I^{\trianglerighteq}(c)$ for a general class of preorders $\trianglerighteq$ that contain all those considered in this paper. The bottomline is that, with Proposition 1, we now have a workable way of computing the critical cost efficiency index.

4.3.1 Proof of Proposition 1

The proof of Proposition 1 is found in full in the Appendix (Section A.1). It uses Theorem 2 in NOQ and here we outline the connection between the two results.

[23] The term 'generalized axiom of revealed preference' as well as its acronym GARP is from the treatment of Afriat's Theorem in Varian (1982). Our concept generalizes the original definition; this is discussed in greater detail in Section 4.3.2.

Given a preorder $\trianglerighteq$, let $\trianglerighteq\bullet\geq$ be the transitive closure of $\trianglerighteq$ and $\geq$.[24] NOQ's result uses the binary relation $\trianglerighteq_e^\bigstar$, where $c^n \trianglerighteq_e^\bigstar c^m$ if there exists $\bar{c} \in I^{\trianglerighteq\bullet\geq}(c^m)$ so that (9) holds. One may state NOQ's Theorem 2 in our context as saying the following: the choices of a subject can be rationalized (at efficiency level e) by a utility function U that agrees with $\trianglerighteq$ if and only if $\trianglerighteq_e^\bigstar$ is strict cycle free in the sense of Definition 8.[25]

Notice that straightforwardly applying NOQ's result would require checking if $c^n \trianglerighteq_e^\bigstar c^m$, which may be problematic because the set $I^{\trianglerighteq\bullet\geq}(c)$ has infinitely many members. The contribution of Proposition 1 is to establish that $\trianglerighteq_e^\bigstar$ *and* $\trianglerighteq_e^*$ *are identical when* $\trianglerighteq$ *is composable*, and $\trianglerighteq_e^*$ is easier to check because $I^{\trianglerighteq}(c)$ is a finite set in many applications. In other words, when the preorder is composable, NOQ's Theorem 2 becomes fully operational as a revealed preference test.[26]

4.3.2 Afriat's Theorem

Proposition 1 delivers a version of Afriat's Theorem as a special case. If one chooses $\trianglerighteq$ as the trivial equality preorder (by definition, $c \trianglerighteq \bar{c}$ if $c = \bar{c}$) then the revealed preference relation in Definition 7 reduces to $c^n \trianglerighteq_e^* c^m$ if either $ep^n \cdot c^n \geq p^n \cdot c^m$ or $c^n \geq c^m$, and the acyclicality condition in Definition 8 reduces to GARP (at efficiency level e) as it is usually defined. Thus Proposition 1 with the equality preorder recovers the known result that the critical cost efficiency index $\bar{e}$ (as defined by (4)) is equal to the supremum of those efficiency levels at which GARP holds.[27]

4.4 Functional Form Restrictions on Utility

We have already observed that when $U : \mathbb{R}_+^4 \to \mathbb{R}$ has the DEU form, both the lottery equivalence and impatience properties are satisfied. More generally, lottery equivalence holds if U is *weakly separable* across states, i.e., there is a strictly increasing function $f : \mathbb{R}_+^2 \to \mathbb{R}$ and a strictly increasing and symmetric function $G : \mathcal{R}(f) \to \mathbb{R}$ (where $\mathcal{R}(f)$ denotes the range of f) such that

$$U((a, b), (a', b')) = G(f(a, b), f(a', b')) \tag{10}$$

[24]Let $a \mathrel{\bar{\trianglerighteq}} b$ if either $a \trianglerighteq b$ or $a \geq b$. The transitive closure of $\trianglerighteq$ and $\geq$ is the binary relation $\trianglerighteq^+$, where $c \trianglerighteq^+ c'$ if there is a sequence $c^1, \ldots, c^K$ so that $c \mathrel{\bar{\trianglerighteq}} c^1 \mathrel{\bar{\trianglerighteq}} c^2 \mathrel{\bar{\trianglerighteq}} \ldots \mathrel{\bar{\trianglerighteq}} c^K \mathrel{\bar{\trianglerighteq}} c'$.

[25]A nuance is that the strict relations $\triangleright_e^*$ and $\triangleright_e^\bigstar$ are defined slightly differently. This is explained in the proof of Lemma 2 in the Appendix.

[26]Polisson, Quah, and Renou (2017) use Theorem 2 in NOQ to test for symmetry and other properties of the utility function. The algorithm they employ can be thought of as an application of Proposition 1, in the sense that the preorders they impose are all composable.

[27]For a direct proof of this claim see Halevy et al. (2018).

This is known as a weakly separable form because for an agent with such a utility, the agent's preference between $((a, b), c_2)$ and $((\bar{a}, \bar{b}), c_2))$ hinges only on whether $f(a, b)$ is smaller or greater than $f(\bar{a}, \bar{b})$ and is independent of the consumption stream (c_2) in state 2. Hence, for an agent with such a preference, it makes sense to speak of the agent's preference over the consumption stream realized in state s, without reference to consumption in other states. We shall refer to a utility function taking the form in (10) as **state-separable**.

Note that state-separable is *not* implied by lottery equivalence. For example,

$$U((a, b), (a', b')) = f(a, b) + h(a, b)h(a', b') + f(a', b') \tag{11}$$

obeys lottery equivalence but cannot be expressible in the form (10).[28] Lottery equivalence also holds if U is separable across time, i.e., there are strictly increasing functions H and g, with g symmetric, such that

$$U((a, b), (a', b')) = H(g(a, a'), g(b, b')) \tag{12}$$

For the functional form (10), impatience holds if f has the property that $f(a, b) \geq f(b, a)$ whenever $a \geq b$. For example, impatience holds if $f(a, b) = u(a) + \delta u(b)$ where u is strictly increasing and continuous, and $\delta \in (0, 1]$. In the case of (11), impatience is satisfied if h (and f) has the property that $h(a, b) \geq h(b, a)$ whenever $a \geq b$.

An important special case of (10) is the *Kihlstrom-Mirman* form (after the representation in Kihlstrom and Mirman (1974)), where there are strictly increasing functions ϕ and u and a number $\delta > 0$ such that

$$U(c) = \phi\Big(u(c_{1,1}) + \delta u(c_{1,2})\Big) + \phi\Big(u(c_{2,1}) + \delta u(c_{2,2})\Big) \tag{13}$$

This form obeys lottery equivalence and, if $\delta \in (0, 1]$, impatience as well.

5 Experimental Results

In this section, we use the NOQ method (as captured by Proposition 1) to test different restrictions on the utility function against the experimental data. We consider several different models, but all of them are weaker than discounted expected utility. By looking at the variation in their performances, we aim to establish precisely which features of the DEU model are less supported by observed choices.

[28]Notice that the ranking between two state 1 consumption streams, say (a, b) and $(\bar{a}, \bar{b})$, clearly depends on the state 2 consumption stream (a', b').

We begin by addressing the first and most elementary issue: are the subjects maximizing *any* utility function?

5.1 Testing For Utility Maximization

Table 2: Pass Rates for Utility Maximization

	$\bar{e} \geq 0.99$	$\bar{e} \geq 0.95$	$\bar{e} \geq 0.90$
Experimental subjects	71.8	90.3	97.1
Uniform random datasets	0.0	0.0	0.0
Resampled datasets	24.3	38.6	65.7

We first test how the participants in the experiments are behaving as imperfect utility-maximizers by calculating their efficiency indices. The basic utility-maximization model appears to perform very well, with virtually everyone having critical cost efficiency indices that exceed 0.9, and more than 70 percent behaving in a way almost indistinguishable from exact utility-maximizers. (See the first row of Table 2.)

To remove the possibility that this is due to an overly lax experimental environment, we use two methods to check the power of our experiment. In both methods, we created 10,000 simulated datasets (each of which has 41 observations) and calculated their efficiency indices. In the first method we generate observations using uniformly random allocations in which the expended tokens sum up to 100. In the second method, we first pool the decisions made by *all* subjects in the experiment and then generate datasets (with 41 observations each) by sampling from this set. The results for the first and second method are reported in the second and third rows of Table 2. A vanishingly small number of datasets generated by the first method are consistent with utility-maximization at any reasonable efficiency threshold; for random datasets generated via the second method, the efficiency indices are higher but still not at all close to the results from the (actual) experimental datasets.

Together these results give strong support to the hypothesis that subjects are maximizing some utility function. Let us now move on to test what properties these utility functions are likely to satisfy.

5.2 Three properties of the DEU Model

In this subsection, we focus on three properties of the discounted expected utility model. It is clear that a discounted expected utility function satisfies impatience and lottery equivalence. But this utility form has a property stronger than lottery equivalence: it also satisfies *correlation neutrality.*

Definition 9. A utility function $U : \mathbb{R}^4_+ \to \mathbb{R}$ satisfies *correlation neutrality* if

$$U\Big((a,b),(a',b')\Big) = U\Big((a',b),(a,b')\Big) \text{ and} \tag{14}$$
$$U\Big((a,b),(a',b')\Big) = U\Big((a,b'),(a,b)\Big) \tag{15}$$

for all (a,b) and (a',b') in $\mathbb{R}^2_+$.

Notice that equations (14) and (15) together imply lottery equivalence so correlation neutrality is a stronger property than lottery equivalence. Correlation neutrality requires a subject to be indifferent between a lottery that (say) pays out consumption streams of $(0,10)$ and $(5,0)$ with equal probability and another that pays out $(5,10)$ and $(0,0)$ with equal probability: the fact that the second lottery involves correlation, in the sense that both the favorable payouts occur in the same state, does not bother the subject. Plainly, correlation neutrality holds in the DEU model and, more generally, it holds when the utility function is separable across time (see (12)).

Correlation neutrality can also be tested using Proposition 1 via an appropriately chosen preorder (see Appendix). Table 3 reports the performance of four properties (lottery equivalence, impatience, their combination, and lastly correlation neutrality) as measured by their respective efficiency indices. Notice that, for all of these models, the pass rates at a given efficiency threshold is lower than that for (basic) utility-maximization, which is as it should be, since the models tested in Table 2 are all more stringent.

Table 3: Pass Rates for DEU properties

	$\bar{e} \geq 0.99$	$\bar{e} \geq 0.95$	$\bar{e} \geq 0.90$
Lottery Equivalence (LE)	62.1	84.5	93.2
Impatience	65.1	84.5	92.2
LE and Impatience	55.3	79.6	91.3
Correlation Neutrality	14.6	22.3	56.3

5.2.1 Lottery Equivalence

We see from Table 3 that a large proportion of subjects display behavior roughly consistent with the maximization of a lottery equivalent utility function. Imposing lottery equivalence must mean lower pass rates compared to the basic utility-maximization model, but the drop in pass rates are quite modest. For example, while 84.5 percent of our experiment participants have an efficiency index above 0.95 for lottery equivalent utility maximization, the corresponding percentage for utility-maximization (see Table

2) is 90.3 percent. (The efficiency index of an additional six participants fall below 0.95 when lottery equivalence is imposed.)

Lottery equivalence may seem like a very reasonable property, but the imposition of this condition has strong implications over and above that from utility-maximization. In particular, a violation of this property can occur with just a single observation, something that is not possible with utility-maximization (where violations can only be observed with multiple observations). As an example we consider subject 10. Subject 10 has a basic efficiency index of 0.976 but it drops to 0.917 when lottery equivalence is imposed. Why has this occurred? Examining even the subject's very first decision is instructive. In observation 1 the subject faces $p^1 = ((1.25, 2), (1, 0.8))$ and purchases $c^1 = ((48, 0), (40, 0))$. This violates lottery equivalence since the subject could instead have purchased the bundle $((40, 0), (50, 0))$, which any lottery equivalent utility-maximizer would strictly prefer to c^1.

The behavior exhibited by participant 10 is not rare: we find that almost 40 percent of subjects would (at some observation) purchase a bundle which would have been cheaper to buy by merely swapping the states in which consumption was taking place (thus freeing up tokens for more consumption). Of course, these violations, while fairly numerous, usually involve only small cost inefficiencies, which is why the efficiency indices across the experiment population remain high.

We view the fact that 84.5 percent of participants have an efficiency index in excess of 0.95 for lottery equivalence as convincing evidence in favor of the notion that most participants have preferences satisfying this property. A potential objection to this claim would be to assert that our test lacks power to properly reject the notion of lottery equivalent utility maximization. To provide evidence against this claim we introduce a new restriction on utility functions which is equally restrictive as lottery equivalence and compare their performance. A utility function U is said to satisfy *time neutrality* if

$$U((a, b), (a', b')) = U((b, a), (b', a'))$$

for any (a, b) and (a', b'). This property is not an implication of DEU model, unless $\delta = 1$. Since the prices on all four goods are chosen with the same distribution, the correlation neutrality and time neutraility properties must, in formal terms, be equally restrictive. But their performances differ sharply (see Table 4). For example, only 68.9 percent of participants have efficiency indices in excess of 0.95 for the time neutral property, considerably below the 84.5 percent reported for lottery equivalence. To check this rigorously, we apply the one-sided version of the McNemar test (see Fagerland, Lydersen, and Laake (2014)), under the null hypothesis that the probability of a participant having an efficiency index above a certain level for lottery equivalence is the same as that for time neutrality. We report p-values for efficiency thresholds of 0.99, 0.95, and 0.9. We see strong evidence that the pass rates at a given efficiency threshold is different for the two tests.

Table 4: Pass Rates: Lottery Equivalence vs Time Neutrality

	$\bar{e} \geq 0.99$	$\bar{e} \geq 0.95$	$\bar{e} \geq 0.90$
Lottery Equivalence	62.1	84.5	93.2
Time Neutrality	56.3	68.9	79.6
McNemar Test p-values	0.066	0.000	0.000

5.2.2 Correlation Neutrality

It is fair to say that correlation neutrality performs abysmally, with only 22.3 percent of participants having a efficiency index score in excess of 0.95 (see Table 3). Many previous papers have commented on theoretical grounds on the undesirability of this implication of the DEU model. Fishburn (1970), Richard (1975), and Duffie and Epstein (1992) are examples. Our results give a sense of the degree to which this matters: if we wanted to make the modest claim that half of our participants satisfy correlation neutrality then we would have to allow for them to waste as much as 10 percent of their budget in their purchasing decisions; in contrast, the choice behavior of half of the subject population can be *exactly* rationalized by lottery equivalent utility functions.

We explore properties that relax correlation neutrality in Section 5.3.

5.2.3 Impatience

We find that 84.5 percent of subjects have an efficiency index in excess of 0.95 (see Table 3). Notice that this the same proportion that satisfy lottery equivalence at the 0.95 efficiency threshold (although the subjects in the two groups do not coincide exactly).

As in the case of lottery equivalence, there are violations of impatience that could be established with just one observation. For example, subject 36 has a perfect efficiency level for basic utility maximization but the efficiency level drops to 0.9 for impatient utility maximization. As an example of a violation, consider observation 15, where $p^{15} = ((0.5, 1.25), (0.8, 1))$ and the subject purchased $c^{15} = ((100, 0), (0, 50))$. Obviously, the subject could have purchased $((100, 0), (50, 0))$ for only 90 tokens which is a bundle that any impatient utility maximizer must consider at least as good as c^{15}. We find violations of this type to be infrequent; indeed, 81.6 percent of subjects do not have any single-observation violations (such as the one exhibited by subject 36). Overall, the introduction of impatience causes the efficiency index to fall below 0.95 for an additional six subjects.

To provide evidence that our result is not driven by the laxity of the impatience property, we introduce the complementary property of patience. A utility function U satisfies *patience* if $U(c_1, c_2) \geq U(c'_1, c'_2)$ whenever impatience requires $U(c_1, c_2) \leq U(c'_1, c'_2)$ (according to the criteria set out in Definition 6). It is clear that in the context of this experiment, patience and impatience are equally restrictive properties. We find that only 66.0 percent of participants have an efficiency index in excess of 0.95 for the patience property, which is considerably below the 84.4 percent found for the impatience property. The pass rates are reported in Table 5, with the McNemar Test confirming our informal observation on the significance of the difference in pass rates.

Table 5: Pass Rates: Impatience vs Patience

	$\bar{e} \geq 0.99$	$\bar{e} \geq 0.95$	$\bar{e} \geq 0.90$
Impatience	65.1	84.5	92.2
Patience	50.5	66.0	73.8
McNemar Test	0.001	0.000	0.000

5.2.4 Joint Test of Lottery Equivalence and Impatience

We now turn to the hypothesis that each agent has a utility function that satisfies *both* lottery equivalence and impatience.[29] Note that it is quite possible for a subject's behavior (say) to be exactly rationalized by a utility function that obeys impatience and also exactly rationalized by a utility function that obeys lottery equivalence, but for it *not* to be exactly rationalized by a utility function with both properties. More generally, for a given subject, the efficiency index for the combined property can be strictly lower than the subject's efficiency index for lottery equivalence and for impatience separately.

It follows from this observation that, while the pass rates for lottery equivalence and the pass rates for impatience impose upper bounds on the pass rate for the combined property, it is theoretically possible for the performance of the combined property to be much worse than that for the two properties separately. It turns out that that is not the case. The pass rate for the combined property remains high; in particular, almost 80 percent of subjects have efficiency indices exceeding 0.95 for the combined property (see Table 3).

[29] If two properties can be defined as agreement with two distinct preorders, then a utility function will satisfy both properties if and only if it agrees with the transitive closure of the two preorders. So long as this transitive closure is also a well-behaved preorder, testing can be carried out using Proposition 1. See Section 4.2 and the Appendix (Sections A.2 and A.3) for further explanation.

For the remaining tests in this section, we shall maintain lottery equivalence and impatience as background assumptions. Not only are these properties ubiquitously assumed in the theory and also intuitive, our experimental evidence has confirmed their viability as maintained assumptions.

5.2.5 Testing DEU

So far we have discussed our findings on lottery equivalence, correlation neutrality, and impatience. While these properties are implied by the DEU model, they are not sufficient to guarantee the DEU form. For the DEU model to hold, the independence axiom would have to be added to these properties.[30] However, the independence axiom cannot be tested *in isolation* using Proposition 1, because there is no preorder $\trianglerighteq$ that represents the independence axiom. Nonetheless, it *is* possible to test the independence axiom jointly with the other properties, since there is a nonparametric test for the DEU model (specifically, to test if U has the form given by (6)) using the method in Polisson et al. (2017). We did not perform that test because it is clear that the model would perform poorly: its pass rates at any efficiency threshold would have to be even lower than that for correlation neutrality.

5.3 Correlation Attitudes

We know that correlation neutrality is broadly rejected by the data. We now investigate whether there is anything more that we can say about correlation attitudes and, in particular, whether there is support for the correlation aversion property.

Definition 10. Let $U : \mathbb{R}^4_+ \to \mathbb{R}$ be a utility function satisfying lottery equivalence. U satisfies *correlation aversion* if for all payouts $a \geq a'$ and $b \geq b'$

$$U\Big((a', b), (a, b')\Big) \geq U\Big((a, b), (a', b')\Big) \qquad (16)$$

[30]To be specific, we would need the following version (due to Debreu (1959)) of the independence axiom which is appropriate in our context:

$$(c_1', c_2) \succeq (c_1, c_2') \text{ and } (c_1'', c_2') \succeq (c_1', c_2'') \implies (c_1'', c_2) \succeq (c_1, c_2''), \quad \forall c_1, c_1', c_1'', c_2, c_2', c_2'' \in \mathbb{R}^2_+$$

If $\succeq$ satisfies this property and correlation neutrality, then Theorem 11.1 in Fishburn (1970) shows that $\succeq$ can be represented by $U(c_{1,1}, c_{1,2}, c_{2,1}, c_{2,2}) = u_1(c_{1,1}) + u_2(c_{1,2}) + u_1(c_{2,1}) + u_2(c_{2,2})$ for increasing and continuous functions u_1 and u_2. From this, it is not difficult to show that $u_2 = \delta u_1$ for $\delta \in (0, 1]$ if impatience and the following property holds:

$$((a, d), (a', d)) \succeq ((b, d), (b', d)) \iff ((d, a), (d, a')) \succeq ((d, b), (d, b')) \quad \forall a, a', b, b', d \in \mathbb{R}_+.$$

If the above inequality is reversed then U is *correlation seeking*.

Since we assume that U satisfies lottery equivalence, if (16) holds then so does $U((a,b'),(a',b)) \geq U((a,b),(a',b'))$. Whether it is this inequality or (16), a subject who satisfies correlation aversion prefers a lottery where high and low payouts occur at each state, over a lottery where all the high payouts are concentrated in one state.

Obviously the DEU model is consistent with correlation aversion since it implies the even stronger property of correlation neutrality. In the case of the Kihlstrom-Mirman model (see (13)), we know from the analysis of Richard (1975) that if ϕ is concave then U displays correlation aversion.

Table 6 reports the results of our tests of correlation aversion and correlation seeking. Both of these properties are combined with lottery equivalence and impatience in our tests, so the pass rates obtained must be lower than those reported in the third row of Table 3. Note that we are not testing correlation seeking because we think it is a plausible global property for a subject's preference; instead, we are testing this property in order to confirm the power of our test to reject the correlation aversion property.[31]

Notice that 75.7 percent of subjects have efficiency indices of 0.95 or higher; in other words, 75.7 percent of subjects are e-rationalized (for *any* $e \in [0, 0.95)$) by utility functions that obey (jointly) lottery equivalence, impatience and correlation aversion. This is only marginally lower than the proportion (79.6 percent) of those who, at the same efficiency threshold, satisfy lottery equivalence and impatience jointly (see Table 3). Furthermore, in our experimental environment, the correlation seeking hypothesis is formally no more or less restrictive than correlation aversion, and yet it is clear from Table 6 that the pass rate for this feature is significantly lower. Together, these results give strong support to the correlation aversion property.

Table 6: Pass Rates: Correlation Aversion vs Correlation Seeking

	$\bar{e} \geq 0.99$	$\bar{e} \geq 0.95$	$\bar{e} \geq 0.90$
Correlation Aversion	51.5	75.7	89.3
Coreelation Seeking	16.5	26.2	52.4
McNemar Test	0.000	0.000	0.000

Another way of comparing the performance of correlation aversion and correlation seeking is to consider the proportion of subjects who have a correlation aversion efficiency level which is higher than their correlation seeking efficiency level. Table 7 shows that 67.96 percent of participants have a higher efficiency level for correlation aversion

[31]We used the same idea in Section 5.2.1 when testing lottery equivalence and time neutrality.

compared to correlation seeking while only 5.83 percent of participants have a higher efficiency level for correlation seeking compared to correlation aversion.

Table 7: Efficiency Levels: Correlation Aversion vs Correlation Seeking

aversion > seeking	seeking > aversion	tie
67.96 %	5.83 %	26.21 %

The support we find in our data for correlation aversion as opposed to correlation neutrality suggests that the behavior of our experimental subjects is not well represented by the DEU model, and more generally, by any model that is separable across time (see (12)). On the other hand, there is some experimental evidence indicating that a subject's preference for lotteries that pay out at time t are unaffected by the presence of lotteries, with independent realizations, in a future period t', which supports a time-separable utility formulation (see Andreoni et al. (2017)). In fact, it is possible to reconcile these results because there are utility functions that satisfy correlation aversion and are also time-separable when lotteries are independent. We discuss this in greater detail in the Appendix (Section A.5).

5.4 Stochastic Impatience

The concept of stochastic impatience has been introduced recently in DeJarnette et al. (2019). Consider the following two lotteries. Lottery C pays out, with equal probability, either 100 dollars in a day's time or 20 dollars in a month. Lottery D pays out, again with equal probability, 20 dollars in a day's time or 100 dollars in one month. It seems reasonable that an 'impatient' person might prefer lottery C to lottery D: in both lotteries, the realized payouts are the same, but lottery C delivers the higher payout sooner. More generally, this preference is captured by the following definition.

Definition 11. A utility function $U : \mathbb{R}^4_+ \to \mathbb{R}$ satisfies *stochastic impatience*[32] if it satisfies lottery equivalence and for all $d \leq b \leq a$,

$$U\Big((a,d),(d,b)\Big) \geq U\Big((b,d),(d,a)\Big). \tag{17}$$

If the above inequality is reversed we say that U satisfies stochastic patience.

Note that stochastic impatience is *not* implied by impatience as we have defined it, even though it *is* implied by correlation neutrality and impatience jointly and thus,

[32] This definition, applicable to consumption streams, is in the Appendix of DeJarnette et al. (2019).

importantly, it is a feature of the DEU model. Indeed,

$$U((b,d),(d,a)) = U((d,d),(b,a)) \leq U((d,d),(a,b)) = U((a,d),(d,b))$$

where the first and last equalities follow from the lottery equivalence and the inequality from impatience. In the case of the Kihlstrom-Mirman model, we show that stochastic impatience holds if (and essentially only if) $u(r) \geq 0$ for all $r \geq 0$ and ϕ has the property that the function Φ given by

$$\Phi(z) = \phi(\exp(z)) \tag{18}$$

is a convex function (see Proposition 3 in the Appendix).[33,34]

Table 8 presents the results of our test of stochastic impatience (both with and without correlation aversion imposed). We see that 79.6 percent of subjects possess efficiency indices in excess of 0.95 for stochastic impatience when correlation aversion is not imposed. Given that we are in fact testing for the combined properties of lottery equivalence, impatience and stochastic impatience, this result is as strong as it could be since (see Table 3) the pass rate of lottery equivalence and impatience is also 79.6 percent. In other words, every subject in our experiment that can be 0.95-rationalized by a utility function that obeys lottery equivalence and impatience can also be 0.95-rationalized by a utility function that obeys lottery equivalence, impatience and stochastic impatience. A similar result holds when correlation aversion is imposed. We see that 75.7 percent of subjects have efficiency indices above 0.95 when correlation aversion is imposed in addition to stochastic impatience. This is as well as the test can perform given the results on correlation aversion in Table 6.

We can also compare the performance of stochastic impatience with the opposite property of stochastic patience. We emphasize that we do *not* consider stochastic patience to be a plausible global property for a subject's preference, but if the stochastic impatience test was too lax or underpowered then we should expect the stochastic patience test to yield similar results. The pass rates for stochastic patience are reported in Table 8 along with the McNemar Test's p-values. We see that there is some evidence that the test has lower power at the the 90% efficiency level; however, there is a statistically significant difference in performance (at the 5% level) at 99% and 95% efficiency.

We next compare the efficiency levels for stochastic impatience and stochastic pa-

[33]When ϕ is twice differentiable, this is equivalent to ϕ having a coefficient of relative risk aversion no higher than 1.

[34]A similar proposition appears in Dillenberger et al. (2018) where a slightly different definition of stochastic impatience is used.

Table 8: Pass Rates: Stochastic Impatience vs Stochastic Patience

	no corr aversion			corr aversion		
	$\bar{e} \geq 0.99$	$\bar{e} \geq 0.95$	$\bar{e} \geq 0.90$	$\bar{e} \geq 0.99$	$\bar{e} \geq 0.95$	$\bar{e} \geq 0.90$
Stochastic Impatience	53.4	79.6	91.3	50.5	75.7	90.3
Stochastic Patience	47.6	75.7	90.3	42.7	70.9	88.3
McNemar Test	0.017	0.023	0.159	0.002	0.013	0.159

tience (see Table 9). There is a large block of participants for whom the revealed preference test is unable to distinguish as being more in line with stochastic impatience or stochastic patience, which indicates an issue with the power of the test. However, in those cases where the efficiency levels differ, a large majority behave more in line with stochastic impatience.

Table 9: Efficiency Levels: Stochastic Impatience vs Stochastic Patience

	impatience > patience	patience > impatience	tie
no corr aversion	13.59 %	2.91 %	83.50 %
corr aversion	14.56 %	1.94 %	83.50 %

Theorem 1 in DeJarnette et al. (2019) shows that if preferences satisfy the independence axiom and several other (more innocuous) properties then stochastic impatience implies RSTL (see Section 2 for a discussion of the Risk Seeking over Time Lotteries (RSTL) concept). Given the evidence they find against RSTL they suggest either relaxing independence or relaxing stochastic impatience.[35] As we have seen, a significant number our subjects behave in accordance with stochastic impatience, which suggests that relaxing the independence axiom may be the preferred option for these participants.

5.5 Takeaways and Challenges

A sharp picture emerges from the tests in this section. Subjects broadly behave as though they are maximizing a lottery equivalent and impatient utility function. There is very little evidence for correlation neutrality, which is a key implication of discounted expected utility. On the other hand, correlation aversion appears to be a widely-held property. We also find that the data is broadly consistent with stochastic impatience, in the sense that most subjects behave in a way that does not contradict this property, though in this case there could be an issue with the power of our test.[36] Overall, as many as three-quarters of subjects could be classified as maximizing (at the 95% efficiency

[35]Our experiment does not allow us to test for RSTL as the property has no testable implications when there are only two time periods.

[36]More specifically, in the power of our experimental environment to test this property.

level) a utility function that is lottery equivalent, impatient, correlation averse, and stochastically impatient (see Table 8).

The evidence put forth provides useful guidance to researchers working on models of choice over time and risk. The ideal model should allow for (not necessarily *require*) correlation aversion and stochastic impatience; it should also be consistent with violations of RSTL (see DeJarnette et al. (2019)) and perhaps permit some limited degree of time separability (see Andreoni et al. (2017)).

6 Ordinal Dominance and a Parametric Model

In this section we estimate a parameterization of the Kihlstrom-Mirman model of equation (13).[37] Recall that the Kihlstrom-Mirman model implies that preferences are state separable and so we first endeavor to test for this property.

6.1 Ordinal Dominance

Suppose we learn that a person prefers to receive 2 dollars today and 2 dollars tomorrow to receiving 3 dollars today. Can we say anything about this person's preference over consumption streams when risk is involved? Perhaps it is reasonable to expect that the person's preference will be unchanged if these two consumption streams are mixed with a third stream; in our notation $((2,2),(a'',b''))$ should be preferred to $((3,0),(a'',b''))$ for any $a'',\ b'' \geq 0$. The notion that a person's preference over lotteries of consumption streams ought to respect the same person's preferences over non-stochastic consumption streams is the basic idea behind the concept of ordinal dominance as defined in Chew and Epstein (1990).

Definition 12. A utility function $U : \mathbb{R}^4_+ \to \mathbb{R}$ satisfies *ordinal dominance* if U satisfies lottery equivalence and

$$U\Big((a,b),(a,b)\Big) \geq U\Big((a',b'),(a',b')\Big) \implies U\Big((a,b),(a'',b'')\Big) \geq U\Big((a',b'),(a'',b'')\Big)$$

for all $a, a', a'', b, b', b'' \in \mathbb{R}_+$.

[37] Note that the Kihlstrom-Mirman model, while useful for estimating separate coefficients for risk and time preferences, has the drawback that when it satisfies stochastic impatience it also satisfies the property of risk seeking over time lotteries, a property challenged by experimental findings in DeJarnette et al. (2019). Notwithstanding this evidence, we still believe the estimation of the parameters of the Kihlstrom-Mirman model a useful exercise as it allows for comparison to previous studies in which measures of risk aversion and intertemporal substitution have been estimated.

Notice that any utility function U also induces a preference over non-stochastic consumption streams. This preference is represented by the increasing and continuous function $f : \mathbb{R}^2_+ \to \mathbb{R}$, given by

$$f(a, b) = U((a, b), (a, b)) \tag{19}$$

Ordinal dominance is equivalent to U being state-separable, i.e., has the form given by (10), with the sub-utility over consumption streams given by f (as defined by (19)). (See Proposition 4 in the Appendix.)

The NOQ method does not allow us to test weak separability as such, essentially because one could not define an underlying preorder on $\mathbb{R}^4_+$ for which consistency with that order is equivalent to weak separability. However, we *can* check for a simple but striking implication of weak separability.

Suppose an agent's utility is weakly separable across states, with the form (10), and that $\hat{c} = (\hat{c}_1, \hat{c}_2)$ maximizes utility at the price vector $\hat{p} = (\hat{p}_1, \hat{p}_2)$. Then

$$\hat{c}_1 \text{ maximizes } f(c_1) \text{ in the set } \{c_1 \in \mathbb{R}^2_+ : \hat{p}_1 \cdot c_1 \leq \hat{p}_1 \cdot \hat{c}_1\}.$$

In other words, if $(\hat{c}_1, \hat{c}_2)$ maximizes overall utility, then $\hat{c}_1$ must maximize the sub-utility in state 1, among consumption streams that cost no more than $\hat{c}_1$: if this were false then an alternative consumption stream $\tilde{c}_1$ is available that raises state 1 sub-utility while not affecting state 2 sub-utility (since $\tilde{c}_1$ costs no more than $\hat{c}_1$) and hence $(\hat{c}_1, \hat{c}_2)$ is not overall optimal. By the same argument, $\hat{c}_2$ maximizes state 2 sub-utility, among all consumption streams in state 2 that cost no more than $\hat{c}_2$.

This observation guarantees that a necessary (but not sufficient) condition for $\mathcal{O}$ to be rationalized by a utility function that satisfies impatience and is weakly separable across states is that the dataset $\mathcal{O}_{\text{split}}$, given by

$$\mathcal{O}_{\text{split}} = \left\{(c_1^1, p_1^1), \ldots, (c_1^N, p_1^N)\right\} \cup \left\{(c_2^1, p_2^1), \ldots, (c_2^N, p_2^N)\right\} \tag{20}$$

can be rationalized by some strictly increasing continuous function $f : \mathbb{R}^2_+ \to \mathbb{R}$ such that $f(a, b) \geq f(b, a)$ if $a \geq b$. ($\mathcal{O}_{\text{split}}$ is the two-commodity dataset formed by splitting each observation in $\mathcal{O}$ into two and combining them into a new dataset.) The efficiency index for this property can be straightforwardly calculated using the NOQ method.

We find that 78.6 percent of participants have a efficiency index of at least 0.95. (See the first row of Table 10.) Since this is based on a test on the reconstructed dataset $\mathcal{O}_{\text{split}}$, it does not nest the tests of lottery equivalence and impatience, which are tested on $\mathcal{O}$. To make the numbers comparable we report on the percent of subjects who have an efficiency index for ordinal dominance above 0.95 *and* have an efficiency index for lottery equivalence with impatience also above 0.95. So, for example, 72.8 percent of the subjects pass both necessary tests of ordinal dominance at the 0.95 efficiency threshold. Overall, these numbers are at least suggestive that a model with ordinal dominance can describe well the behavior we observe.

Table 10: Ordinal Dominance Test

	$\bar{e} \geq 0.99$	$\bar{e} \geq 0.95$	$\bar{e} \geq 0.90$
Impatient subutility	65.0	78.6	84.5
Impatient subutility & LE+Impatience	46.6	72.8	83.5

6.2 A Parametric and a Semi-parametric model

In this section we estimate two models. First, we estimate a semi-parametric model which uses discounted utility with a power felicity function to evaluate timed payouts within states and a general symmetric aggregator to aggregate utility across states. That is the bundle $c = (c_{1,1}, c_{1,2}, c_{2,1}, c_{2,2})$ is evaluated by some utility function U with form

$$U(c) = G\left(c_{1,1}^{1-\gamma} + \delta c_{1,2}^{1-\gamma}, c_{2,1}^{1-\gamma} + \delta c_{2,2}^{1-\gamma}\right) \tag{21}$$

where G is some increasing, continuous, and symmetric aggregator function and $\gamma \in [0, 1)$.[38] While this model places a lot of structure on how the subject trades off timed payouts within a state it does not place parametric assumptions on how a subject evaluates risky payouts and does not even require that risky payouts be evaluated using expected utility.

Second, we estimate a parametric version of the Kihlstrom-Mirman model (see equation (13)) where we assume that u and ϕ are power functions or ϕ is the log function. That is, we assume that U can be represented either by

$$U(c) = \tfrac{1}{1-\eta}\left[\left(c_{1,1}^{1-\gamma} + \delta c_{1,2}^{1-\gamma}\right)^{\frac{1-\eta}{1-\gamma}} + \left(c_{2,1}^{1-\gamma} + \delta c_{2,2}^{1-\gamma}\right)^{\frac{1-\eta}{1-\gamma}}\right] \tag{22}$$

where $\gamma \in [0, 1)$, $\eta \in [0, \infty)\backslash\{1\}$, and $\delta \in [0, 1]$ or that

$$U(c) = \ln\left(c_{1,1}^{1-\gamma} + \delta c_{1,2}^{1-\gamma}\right) + \ln\left(c_{2,1}^{1-\gamma} + \delta c_{2,2}^{1-\gamma}\right) \tag{23}$$

[38] Our restriction of γ to be less than 1 does not appear to restrict any of our estimates. We make this restriction to accommodate the many corner solutions observed. Note that when $\gamma > 1$ then $0^{1-\gamma}$ is either undefined or $-\infty$.

Since the function ϕ in (13) is concave in (22) and (23), correlation aversion is embedded in our specification (see Richard (1975)). Note that our semi-parametric model nests the parametric model and that the two models coincide for non-risky temporal streams of consumption.

Recall that stochastic impatience holds if (and, in a sense, only if) the function Φ (given by (18)) is a convex function. In this case, if U has the form (23), then $\Phi(z) = z$, and thus stochastic impatience holds; if U has the form (22), then

$$\Phi(z) = \frac{1}{(1-\eta)} \exp\left(\frac{z(1-\eta)}{(1-\gamma)}\right),$$

in which case stochastic impatience is satisfied if $\eta \in (0, 1)$.

The parameters in the model (22) and (23) have natural interpretations. Note that when risk is absent ($\bar{c}_1 = c_{1,1} = c_{2,1}$ and $\bar{c}_1 = c_{1,2} = c_{2,2}$) then preferences over $\bar{c}_1$ and $\bar{c}_2$ have constant elasticity of intertemporal substitution with parameter $1/\gamma$. Thus, γ can be thought of as representing the curvature in the time preference utility function. For the parametric model, when $c_{1,2} = c_{2,2} = 0$ and so we are only considering earlier period consumption then the preferences take the constant relative risk aversion form with relative risk aversion coefficient η (or 1 in the case of (23)). Thus, η represents the curvature of the risk preference utility function.

Table 11: Semi-parametric and Parametric Models

	$\bar{e} \geq 0.99$	$\bar{e} \geq 0.95$	$\bar{e} \geq 0.90$
LE and Impatience	55.3	79.6	91.3
Semi-parametric model	24.3	47.6	73.8
Parametric model	12.6	35.9	59.2

We estimate the models by selecting parameters that maximize the efficiency index. A closely related approach to parametric estimation was first used in Halevy et al. (2018).[39] The distribution of efficiency indices is reported in Table 11. The proportion exceeding a given threshold are considerably lower than the corresponding numbers for the nonparametric test, but this is to be expected given that we now estimating parametric and semi-parametric models. However, note that both parametric and semi-parametric models generally outperform the non-parametric correlation neutral model.

The median annualized discount rate estimated was 0.67 for the parametric model and 0.85 for the semi-parametric model. These estimates are higher than that in An-

[39]The index used by Halevy et al. (2018) is not exactly Afriat's efficiency index but a related measure due to Varian (1990). Halevy et al. (2018) compare the results from their estimation approach with the one obtained by minimizing the sum of square differences and found that their method performs better. The data in that case consists of portfolio decisions involving risk, with no time dimension.

dreoni and Sprenger (2012b) which was 0.41. However, the upper quartile for the estimated discount rate for both models is the complete patience level of 1.0. This may reflect the fact that the difference between early and late payment (8 weeks) was not sufficiently long to pick up meaningful discount rates in some subjects. On the other hand our lower quartile for this parameter is 0.42 for the parametric and 0.35 for the semi-parametric which are close to the aforementioned rate found by Andreoni and Sprenger (2012b).

The median estimated time curvature γ is 0.10 for the parametric model and 0.08 for the semi-parametric model. The lower quartile of our participants have an estimated γ of 0 indicating perfectly linear utility for both models. The top quartile held a moderate value of 0.25 for the parametric model and 0.23 for the semi-parametric model.

We now turn the estimates of the risk parameter η for the parametric model. The median estimate was 0.55. This can be compared to a finding of 0.44 found by Choi et al. (2007) in their experiment run with equiprobable states (but where all payments were made on the same date). Note that Choi et al. (2007) estimates a rank dependent utility model that allows for a kink in the utility function at certainty. This may explain why their estimate of the curvature is lower than ours. Our estimates for the upper and lower quartiles for the η parameter were 1.0 and 0.35 respectively. 78 out of 103 (or 75.7 percent of subjects) have estimated η no greater than 1 and would thus satisfy stochastic impatience.

A striking feature in our estimates for the parametric model is that there is almost *no correlation* between estimated γ and estimated η. In the DEU model these coefficients are identical, which would imply a very high positive correlation coefficient; instead we find that it is -0.043.

Appendix

A.1 Well-behaved preorders and Proposition 1

In the paper we consider preorders defined on the consumption space $\mathbb{R}^4_+$. Formally, a binary relation $\trianglerighteq$ on $\mathbb{R}^4_+$ is a subset of $\mathbb{R}^4_+ \times \mathbb{R}^4_+$. We write $a \trianglerighteq b$ if $(a, b) \in \trianglerighteq$. A *preorder* refers to a binary relation that is reflexive ($a \trianglerighteq a$ for all a) and transitive ($a \trianglerighteq b$ and $b \trianglerighteq c$ implies $a \trianglerighteq c$). A preorder is said to be *closed* if $\trianglerighteq$, considered as a subset of $\mathbb{R}^4_+ \times \mathbb{R}^4_+$, is a closed set. The product order $\geq$ (where $c' \geq c$ if $c'_{ij} \geq c_{ij}$ for $i = 1, 2$, $j = 1, 2$) is a closed preorder.

Given a binary relation $\trianglerighteq$, we can define its *transitive closure*; this is another binary relation on $\mathbb{R}^4_+$, which we shall denote by $\trianglerighteq^+$. By definition, $a \trianglerighteq^+ b$ if there is a finite set of elements c^j (for $j = 1, 2, \ldots, m$) such that $a \trianglerighteq c^1$, $c^1 \trianglerighteq c^2$, and so forth until $c^m \trianglerighteq b$. If $\trianglerighteq$ is a reflexive binary relation, then $\trianglerighteq^+$ is a preorder.

Given two binary relations $\trianglerighteq_1$ and $\trianglerighteq_2$, it its union is also a binary relation; we denote the transitive closure of this union by $\trianglerighteq_1 \bullet \trianglerighteq_2$; we simply refer $\trianglerighteq_1 \bullet \trianglerighteq_2$ as the *transitive closure of* $\trianglerighteq_1$ *and* $\trianglerighteq_2$. Note that if $\trianglerighteq_1$ and $\trianglerighteq_2$ are preorders, its union need not be a preorder, but its transitive closure will be.

Let $\circ$ denote the composition between two preorders $\trianglerighteq_1$ and $\trianglerighteq_2$, so $c \trianglerighteq_1 \circ \trianglerighteq_2 c'$ if there exists $\tilde{c}$ such that $c \trianglerighteq_1 \tilde{c} \trianglerighteq_2 c'$. Obviously, $\trianglerighteq_1 \circ \trianglerighteq_2$ is a binary relation, but it need not be a preorder. A sufficient condition to guarantee that $\trianglerighteq_1 \circ \trianglerighteq_2$ is a preorder is that the two preorders $\trianglerighteq_1$ and $\trianglerighteq_2$ *commute*, i.e., $\trianglerighteq_1 \circ \trianglerighteq_2 = \trianglerighteq_2 \circ \trianglerighteq_1$. Note that a preorder $\trianglerighteq$ is composable in the sense of Definition 4 if and only if it commutes with the $\geq$ preorder.

The following lemma on composable preorders (see Definition 4) has a completely straightforward proof which we shall omit.

Lemma 1. *If* $\trianglerighteq$ *is a closed and composable preorder then* $\geq \circ \trianglerighteq$ *is a closed preorder.*

Definition 13. A preorder $\trianglerighteq$ on $\mathbb{R}^4_+$ is *well-behaved* if it is closed, composable, and $c \trianglerighteq c'$ implies not $c' > c$.

A utility function $U : \mathbb{R}^4_+ \to \mathbb{R}$ is said to *strictly agree* (or *extend*) the preorder $\trianglerighteq$ if $U(c'') \geq U(c')$ whenever $c'' \trianglerighteq c'$ and $U(c'') > U(c')$ whenever $c'' \vartriangleright c'$. Note that this notion is stronger than simply saying that U agrees with $\trianglerighteq$ (see Definition 3). Our proof of Proposition 1 uses the following lemma.

Lemma 2. *Let* $\mathcal{O} = \{(c^1, p^1), \ldots, (c^N, p^N)\}$ *be a dataset, let* $\trianglerighteq$ *be a well-behaved preorder and let* $e < e^*$ *for* $e, e^* \in [0, 1]$*. If* $\mathcal{O}$ *satisfies* e^**-GARP-*$\trianglerighteq$ *then there is a utility function that strictly agrees with* $\trianglerighteq$ *and* e*-rationalizes* $\mathcal{O}$*.*

Proof. Let $e < e^*$ and $B^n(e) = \{c \in \mathbb{R}^4_+ : p^n \cdot c \leq e p^n \cdot c^n\}$. By Definition 2, the utility function U that we seek must satisfy $U(c^n) \geq U(c)$ for all $c \in B^n(e)$. Since U is strictly increasing by definition, this is equivalent to $U(c^n) \geq U(c)$ for all $c \in \bar{B}^n(e)$, where

$$\bar{B}^n(e) = \{c \in \mathbb{R}^4_+ : p^n \cdot c \leq 100e \text{ or } c \leq c^n\}. \tag{24}$$

By Lemma 1, $\trianglerighteq^+ = \geq \circ \trianglerighteq$ is a closed preorder. By Theorem 2 in NOQ, there is a function U that strictly agrees with $\trianglerighteq^+$ and satisfies $U(c^n) \geq U(c)$ for all $c \in \bar{B}^n(e)$ if the choices obey a sort of no-cycling condition. This is precisely what we need provided the following is true: (a) that if U strictly agrees with $\trianglerighteq^+$ then it is strictly increasing (and thus a true utility function in the sense defined in this paper) and strictly agrees with $\trianglerighteq$; (b) the no-cycling condition in NOQ, when specialized to our context, is satisfied.

We first establish (a). Since U strictly agrees with $\trianglerighteq^+$, $U(c'') \geq U(c')$ if $c'' > c'$ or $c'' \triangleright c'$, and $U(c'') = U(c')$ only if $c' \trianglerighteq^+ c''$ (here $c'' \triangleright c'$ means $c'' \trianglerighteq c'$ and not $c' \trianglerighteq c''$). The latter condition means there is a $\bar{c}$ so that $c' \geq \bar{c} \trianglerighteq c''$. Note that $c' \geq \bar{c} \trianglerighteq c''$ means not $c'' > c'$ (by $\trianglerighteq$ being well-behaved). Now, suppose $c'' \triangleright c'$. Then we have $c' \geq \bar{c} \trianglerighteq c'' \triangleright c'$ where the $\geq$ cannot be replaced by $>$ without contradicting the fact that $\trianglerighteq$ is well-behaved and so it must be that $c' = \bar{c} \trianglerighteq c'' \triangleright c'$ which is a contradiction. We conclude that $U(c'') > U(c')$ if $c'' > c'$ or $c'' \triangleright c'$.

We now turn to claim (b). NOQ's condition is that e-GARP-$\trianglerighteq$ holds except that the revealed preference relations in NOQ are defined differently. We need to show that e^*-GARP-$\trianglerighteq$ guarantees that NOQ's e-GARP-$\trianglerighteq$ holds. In NOQ, c^n is said to be e-revealed preferred to c^m if there is $\tilde{c}$ such that $\tilde{c} \in \bar{B}^n(e)$ and $\tilde{c} \trianglerighteq^+ c^m$, and c^n is strictly e-revealed preferred to c^m if $\tilde{c}$ can be chosen such that $\tilde{c} \triangleright^+ c^m$. Let us denote NOQ's revealed preference relations by $\trianglerighteq_e^\bigstar$ and $\triangleright_e^\bigstar$.

It is easy to see that $c^n \trianglerighteq_{e^*}^* c^m$ if and only if $c^n \trianglerighteq_{e^*}^\bigstar c^m$ and so $c^n \trianglerighteq_e^\bigstar c^m$ implies $c^n \trianglerighteq_{e^*}^* c^m$. It can also be shown easily that $c^n \triangleright_{e^*}^* c^m$ implies $c^n \triangleright_{e^*}^\bigstar c^m$. However, $c^n \triangleright_{e^*}^\bigstar c^m$ does not imply $c^n \triangleright_{e^*}^* c^m$. Yet, it is the case that $c^n \triangleright_e^\bigstar c^m$ implies $c^n \triangleright_{e^*}^* c^m$ which will be used to establish (b). Let us prove this claim. Suppose there exists $\tilde{c} \in \bar{B}^n(e)$ so that $\tilde{c} \triangleright^+ c^m$. By the composable property of $\trianglerighteq$ this means there exists $\bar{c} \trianglerighteq c^m$ so that $\bar{c} \in \bar{B}^n(e)$. This is equivalent to $\bar{c} \trianglerighteq c^m$ and

$$ep^n \cdot c^n \geq p^n \cdot \bar{c} \qquad \text{or} \qquad c^n \geq \bar{c}$$

Since $e^* > e$, we obtain $c^n \triangleright_{e^*}^* c^m$. Now, suppose the data violate e^*-GARP-$\trianglerighteq$ with the $\trianglerighteq_{e^*}^*$ revealed preference relation. From what we have established about $\trianglerighteq_e^\bigstar$, it follows that e-GARP-$\trianglerighteq$ is violated for the $\trianglerighteq_e^\bigstar$ revealed preference relation. Thus, e^*-GARP-$\trianglerighteq$ in the sense of Definition 8 implies e-GARP-$\trianglerighteq$ in NOQ's sense. □

Proof of Proposition 1. We first show that if the dataset $\mathcal{O}$ can be rationalized by a utility function U that agrees with $\trianglerighteq$ then the data will satisfy e-GARP-$\trianglerighteq$. So, let U be a utility function that agrees with $\trianglerighteq$ and rationalizes $\mathcal{O}$ at efficiency level e. Define $\bar{B}^n(e)$ by (24). If $c^n \trianglerighteq_e^* c^m$, then by definition there is $\bar{c} \in \bar{B}^n(e)$ such that $\bar{c} \trianglerighteq c^m$. This implies that $U(c^n) \geq U(\bar{c}) \geq U(c^m)$. If $c^n \triangleright_e^* c^m$ then we have, in addition, that $U(c^n) > U(\bar{c})$. It follows from this that e-GARP-$\trianglerighteq$ holds. Thus

$$\sup\Big\{e \in [0,1] : \mathcal{O} \text{ satisfies } e\text{-GARP-} \trianglerighteq \Big\} \geq \sup\Big\{e \in [0,1] : \mathcal{O} \text{ is } e\text{-rationalizable-} \trianglerighteq \Big\}$$

We must now establish the "$\leq$" inequality part of the proof. Suppose e^* is the supremum value of all e so that $\mathcal{O}$ satisfies e-GARP-$\trianglerighteq$. Lemma 2 shows that for any $e < e^*$ we may find a utility function U that strictly agrees with $\trianglerighteq$ and rationalizes $\mathcal{O}$ at efficiency level e. As any U that strictly agrees with $\trianglerighteq$ also agrees with $\trianglerighteq$, the supremum over all e such that $\mathcal{O}$ is e-rationalizable-$\trianglerighteq$ must be weakly greater than e^*. This establishes the "$\leq$" part of the proof. $\square$

A.2 Testing Properties Jointly

Let $\trianglerighteq_1 \bullet \trianglerighteq_2$ be the transitive closure of the preorders $\trianglerighteq_1$ and $\trianglerighteq_2$. Suppose we would like to know if there is a utility function U that agrees with both $\trianglerighteq_1$ and $\trianglerighteq_2$ and rationalizes the choices of a participant at efficiency level e. This is clearly equivalent to asking if the choices are e-rationalizable-($\trianglerighteq_1 \bullet \trianglerighteq_2$). The following property, possessed by all preorders considered in this paper, makes it possible to test this condition.

Definition 14. A preorder $\trianglerighteq$ is an arranging preorder if $(\bar{c}_1, \bar{c}_2, \bar{c}_3, \bar{c}_4) \trianglerighteq (c_1, c_2, c_3, c_4)$ implies $\bar{c}_k \in \{c_1, c_2, c_3, c_4\}$ for each k.

Lemma 3. *If $\trianglerighteq_1$ and $\trianglerighteq_2$ are composable and arranging preorders, then $\trianglerighteq_1 \bullet \trianglerighteq_2$ is a composable and arranging preorder.*

Proof. Clearly $\trianglerighteq_1 \bullet \trianglerighteq_2$ is arranging. Suppose $c \geq \circ \; (\trianglerighteq_1 \bullet \trianglerighteq_2) \; c'$. Then it must be the case that $c \geq \circ \trianglerighteq_1 \circ \trianglerighteq_2 \circ \trianglerighteq_1 \circ \trianglerighteq_2 \ldots \trianglerighteq_1 \circ \trianglerighteq_2 c'$. By applying the composable property repeatedly, we see $c \; (\trianglerighteq_1 \bullet \trianglerighteq_2) \circ \geq c'$. $\square$

Proposition 2. *Suppose $\trianglerighteq_1$ and $\trianglerighteq_2$ are well-behaved and arranging preorders and suppose the following property holds: $c \; (\trianglerighteq_1 \bullet \trianglerighteq_2) \; c'$ implies not $c' > c$. Then $\trianglerighteq_1 \bullet \trianglerighteq_2$ is a well-behaved and arranging preorder.*

Proof. By Lemma 3 we see that $\trianglerighteq_1 \bullet \trianglerighteq_2$ is composable and arranging. So, to show that $\trianglerighteq_1 \bullet \trianglerighteq_2$ is well-behaved we must show that it is closed.

For a binary relation $\trianglerighteq$ and integer K let $\trianglerighteq^K$ denote the binary relation formed by composing $\trianglerighteq$ with itself K times. As $\trianglerighteq_1$ and $\trianglerighteq_2$ are arranging it must be the case that there is a K so that $\trianglerighteq_1 \bullet \trianglerighteq_2 = \left(\trianglerighteq_2 \circ \trianglerighteq_1 \right)^K$. Thus, it suffices to show for two closed and arranging preorders $\trianglerighteq_1$ and $\trianglerighteq_2$ that $\trianglerighteq_2 \circ \trianglerighteq_1$ is closed. Suppose we have a sequence $c_k \to c$ and $c'_k \to c'$ so that $c_k \; \trianglerighteq_2 \circ \trianglerighteq_1 \; c'_k$ for each k. So, there exists a sequence $\bar{c}_k$ so that $c_k \trianglerighteq_2 \bar{c}_k \trianglerighteq_1 c'_k$. We may assume that $\bar{c}_k \to \bar{c}$ (or at least we could create such a subsequence). As $\trianglerighteq_1$ and $\trianglerighteq_2$ are closed we see $c \trianglerighteq_2 \bar{c} \trianglerighteq_1 c'$. $\square$

In order to check if a dataset is consistent with two properties (such as lottery equivalence and impatience) jointly, we would need to apply Proposition 1 to the preorder that is the transitive closure of two preorders. In this situation, Proposition 2 is useful because it spells out the condition under which the transitive closure of well-behaved and arranging preorders is also well-behaved and arranging.

A.3 Our menagerie of preorders

Except for the elementary results relying on Afriat's Theorem in Section 5.1, all the results in Section 5 rely on applications of Proposition 1, with different preorders chosen for different tests. We know that a utility function satisfies lottery equivalence if and only if it agrees with $\trianglerighteq_{LE}$ and impatience if and only if it agrees with $\trianglerighteq_I$. Given these, there are obvious analogous preorders $\trianglerighteq_{TN}$ and $\trianglerighteq_P$ that characterize time neutrality and patience. Other properties studied in Section 5 can be characterized in a similar way. The following claims are easy to check:
(1) A utility function satisfies correlation neutrality if and only if it agrees with $\trianglerighteq_{CN}$, where $\trianglerighteq_{CN}$ is defined as follows: for all a, b, a', and $b' \geq 0$, we have $((a,b),(a',b')) \trianglerighteq_{CN}$ $((a',b),(a,b'))$, $((a,b),(a',b')) \trianglerighteq_{CN}$ $(a,b'),(a',b))$, and $((a,b),(a',b')) \trianglerighteq_{CN}$ $((a',b'),(a,b))$.
(2) A utility function satisfies correlation aversion if and only if it agrees with $\trianglerighteq_{LE}$ and with $\trianglerighteq_{CA}$, where $\trianglerighteq_{CA}$ is defined as follows: if $a \geq a'$ and $b \geq b'$, then $((a',b),(a,b')) \trianglerighteq_{CA}$ $((a,b),(a',b'))$. The correlation loving property can be characterized by requiring the utility function to agree with $\trianglerighteq_{LE}$ and with $\trianglerighteq_{CL}$, where the latter is defined as follows: $c' \trianglerighteq_{CL} c''$ if $c'' \trianglerighteq_{CA} c'$.
(3) A utility function satisfies stochastic impatience if and only if it agrees with $\trianglerighteq_{LE}$ and with $\trianglerighteq_{SI}$, where $\trianglerighteq_{SI}$ is defined as follows: for all d', d'', $d \leq b \leq a$ where $d = \max(d', d'')$, we have $((a,d),(d,b)) \trianglerighteq_{SI}$ $((b,d'),(d'',a))$. A utility function satisfies stochastic patience if it agrees with $\trianglerighteq_{LE}$ and with $\trianglerighteq_{SP}$, with the latter defined as follows: for all d', d'', $d \leq b \leq a$ where $d = \max(d', d'')$, we have $((b,d),(d,a)) \trianglerighteq_{SP}$ $((a,d'),(d'',b))$.

It is straightforward to check that every preorder we consider, i.e., every preorder in the set

$$\{\trianglerighteq_{LE}, \trianglerighteq_I, \trianglerighteq_{TN}, \trianglerighteq_P, \trianglerighteq_{CN}, \trianglerighteq_{CA}, \trianglerighteq_{CL}, \trianglerighteq_{SI}, \trianglerighteq_{SP}\}$$

is well-behaved and arranging. Thus by Proposition 2, the transitive closure of the union of any subset of these preorders is also well-behaved and arranging. For example, in Section 5.3, when we tested correlation aversion, we maintained lottery equivalence and impatience as background assumptions. A utility function will have these properties if and only if it agrees with the preorder $\trianglerighteq \, = \, \trianglerighteq_{LE} \bullet \trianglerighteq_I \bullet \trianglerighteq_{CA}$. This preorder is again well-behaved. Thus, according to Proposition 1, we can calculate the efficiency index for the joint { lottery equivalence + impatience + correlation aversion } property by checking e-GARP according to $\trianglerighteq$.

A.4 Implementing the Tests

Here we discuss how to actually test if the behavior of a subject satisfies e-GARP-$\trianglerighteq$. We shall focus on the case where we seek to test e-GARP-$\trianglerighteq$ for some $\trianglerighteq = \trianglerighteq_1 \bullet \trianglerighteq_2$ however our method shall easily generalize to the case of combining K preorders, $\trianglerighteq = \trianglerighteq_1 \bullet \trianglerighteq_2 \bullet \ldots \bullet \trianglerighteq_K$.

We shall suppose that the preorders $\trianglerighteq_1$ and $\trianglerighteq_2$ are well-behaved and arranging, so that for any c the sets $I^{\trianglerighteq_1}(c)$ and $I^{\trianglerighteq_2}(c)$ are finite and easy to compute.

Here is a procedure which will check for e-GARP-$\trianglerighteq$.

Procedure.

Step 1. *For each m calculate the set $\{c \in \mathbb{R}^4_+ : c \trianglerighteq c^m\}$.*

Step 2. *For each n and m determine if $c^n \trianglerighteq^*_e c^m$.*

Step 3. *Determine if there is a sequence $n_1, \ldots, n_K$ so that $c^{n_1} \trianglerighteq^*_e c^{n_2} \trianglerighteq^*_e \ldots \trianglerighteq^*_e c^{n_K}$ and $c^{n_K} \triangleright^*_e c^{n_1}$*

Step 3 in the Procedure can be accomplished using Warshall's algorithm (an algorithm for computing the transitive closure of a graph.) So, the challenge lies in Steps 1 and 2. To accomplish Step 1 we shall construct a sequence of sets

$$V_0, V_{1,1}, V_{1,2}, V_{2,1}, V_{2,2}, V_{3,1}, V_{3,2}, V_{4,1}, V_{4,2} \ldots$$

whose elements are in $\mathbb{R}^4_+$.[40] Let $V_0 = \{c^m\}$ and let $V_{1,1} = \{c \in \mathbb{R}^4_+ : c \trianglerighteq_1 \bar{c}, \ \text{ where } \bar{c} \in V_0\}$. Let $V_{1,2} = \{c \in \mathbb{R}^4_+ : c \trianglerighteq_2 \bar{c}, \ \text{ where } \bar{c} \in V_{1,1}\}$. And, for $k > 1$ let

$$V_{k,1} = \{c \in \mathbb{R}^4_+ : c \trianglerighteq_1 \bar{c}, \text{ where } \bar{c} \in V_{k-1,2}\}, \quad V_{k,2} = \{c \in \mathbb{R}^4_+ : c \trianglerighteq_2 \bar{c}, \text{ where } \bar{c} \in V_{k,1}\}$$

Each of these sets is computable since $I^{\trianglerighteq_1}(c^m)$ and $I^{\trianglerighteq_2}(c^m)$ can be computed. Because $\trianglerighteq$ is arranging there will be a $k > 1$ so that $V_{k,2} = V_{k-1,2}$. It is easy to show that $V_{k,2} = \{c \in \mathbb{R}^4_+ : c \trianglerighteq c^m\}$ and so we have constructed the desired set.

Once Step 1 is completed Step 2 is easy. Merely, consider each bundle $\bar{c} \in \{c \in \mathbb{R}^4_+ : c \trianglerighteq c^m\}$ and consider if $ep^n \cdot c^n \geq p^n \cdot \bar{c}$. If the answer is "yes" for any $\bar{c}$ then $c^n \trianglerighteq^*_e c^m$ and otherwise it is not the case that $c^n \trianglerighteq^*_e c^m$.

A.5 Separability and Correlation Aversion

Here we discuss the connection between correlation aversion and the separable structures on preferences. Let $X = (X^0, X^1, \ldots, X^T)$ denote a $T+1$ dimensional random vector

[40] The V stands for "visited."

over some finite sample space with S elements. Let us write $X_1, \ldots, X_S$ for the S non-risky temporal streams of consumption. X_s is the consumption stream realized, over $T+1$ periods, should state s occur (so, $X_s \in \mathbb{R}_+^{T+1}$). A utility function U is *state-separable* if X is awarded utility

$$U(X) = F_S\Big(G_S(X_1), \ldots, G_S(X_S)\Big)$$

for some increasing aggregator function F_S and increasing sub-utility function G_S. We say that U is *time-separable* if X is awarded utility

$$U(X) = F_T\Big(G_T(X^0), \ldots, G_T(X^T)\Big)$$

for some increasing aggregator function F_T and increasing sub-utility function G_T (whose domain consists of the random variables X^t). Note that the DEU model is both state-separable and time-separable while a utility function that allows for correlation aversion cannot be time-separable.

A discussion found in Andreoni and Sprenger (2012a), Epper and Fehr-Duda (2015), and Andreoni et al. (2017) concerns whether utility should be time-separable or state-separable. Andreoni et al. (2017) observe that with the time-separable formulation, an agent's preference over lotteries that pay out in period t are unaffected by the presence of lotteries that pay out independently in some other time period $t' > t$; they refer to this property as common future risk invariance (CFRI). Andreoni et al. (2017) find some support for this behavior in their experiment.

It is in fact possible to reconcile correlation aversion with a restricted version of time-separability. For example, suppose that an agent has the model

$$U(X) = -\mathbf{E}\left[\prod_{t=0}^{T} \exp\left(-\delta^t u(X^t)\right)\right]$$

This function[41] is obviously state-separable and it is not difficult to check that it obeys correlation aversion (when states are equiprobable), so it could potentially explain the behavior typically displayed by subjects in our experiment. On the other hand, the model is also *time-separable when* $X^0, \ldots, X^T$ *are independent* (which follows from the fact that $\mathbf{E}[Z_1 Z_2] = \mathbf{E}[Z_1]\mathbf{E}[Z_2]$ when Z_1 and Z_2 are independent random variables), and thus the agent could also display CFRI.

[41] A similar functional form can be found in Richard (1975).

A.6 Stochastic Impatience

Proposition 3. *Suppose that U has the form of* (13)*, where ϕ and u are both strictly increasing and $u(0) = 0$ (so $u(r) > 0$ for all $r > 0$). Then U displays stochastic impatience if Φ as defined by (18) is a convex function. Conversely, if Φ is not a convex function, then there is a strictly increasing and concave function u, with $u(0) = 0$, such that U with the form* (13) *violates stochastic dominance.*

The proof of this result requires the following lemma.

Lemma 4. *Suppose $\phi : \mathbb{R}_{++} \to \mathbb{R}$ is an increasing function. For any $\beta \geq 0$, let $\Phi_\beta(z) = \phi(\exp(z) + \beta)$. Then Φ_β is a convex function if $\Phi(z) = \phi(\exp(z))$ is a convex function of z.*

Proof. It suffices to show that for any $z'' > z'$ and $x > 0$,

$$\Phi_\beta(z'' + x) - \Phi_\beta(z' + x) \geq \Phi_\beta(z'') - \Phi_\beta(z').$$

Note that $\Phi_\beta(z'' + x) - \Phi_\beta(z' + x) = \Phi(\ln(\exp(z'' + x) + \beta)) - \Phi(\ln(\exp(z' + x) + \beta))$. Clearly, $\ln(\exp(z' + x) + \beta)) \geq \ln(\exp(z') + \beta)$ and

$$\begin{aligned} &\ln(\exp(z'' + x) + \beta)) - \ln(\exp(z' + x) + \beta) \\ = &\ln\left(\exp(z'') + \frac{\beta}{\exp(x)}\right) - \ln\left(\exp(z') + \frac{\beta}{\exp(x)}\right) \geq \ln(\exp(z'') + \beta)) - \ln(\exp(z') + \beta), \end{aligned}$$

where the inequality follows from the concavity of the log function. Therefore, we obtain, since Φ is convex and increasing,

$$\Phi(\ln(\exp(z'' + x) + \beta)) - \Phi(\ln(\exp(z' + x) + \beta)) \geq \Phi(\ln(\exp(z'') + \beta)) - \Phi(\ln(\exp(z') + \beta))$$

which is the inequality we want. □

Proof of Proposition 3. Suppose U has the form ((13)) and Φ is convex. We need to show that for all $a \geq b \geq d$

$$U\Big((a, d), (d, b)\Big) \geq U\Big((b, d), (d, a)\Big) \tag{25}$$

This is trivially true if $b = d$, so we assume that $b > d$. Let $\beta = u(d) + \delta u(d)$. Plugging in the representation (13) into (25) we obtain

$$\phi\Big(u(a) - u(d) + \beta\Big) + \phi\Big(\delta(u(b) - u(d)) + \beta\Big) \geq \phi\Big(u(b) - u(d) + \beta\Big) + \phi\Big(\delta(u(a) - u(d)) + \beta\Big) \tag{26}$$

Equivalently,

$$\Phi_\beta\Big(\ln(u(a)-u(d))\Big)-\Phi_\beta\Big(\ln(\delta)+\ln(u(a)-u(d))\Big)\geq$$
$$\Phi_\beta\Big(\ln(u(b)-u(d))\Big)-\Phi_\beta\Big(\ln(\delta)+\ln(u(b)-u(d))\Big) \quad (27)$$

This clearly holds if Φ_β is convex, which it is by Lemma 4 since $\beta \geq 0$.

If Φ is not convex, then there is $v'' > v'$ and $w > 0$ such that

$$\Phi(v'')-\Phi(v''-w)<\Phi(v')-\Phi(v'-w).$$

Re-arranging this gives

$$\phi(\exp(v''))+\phi(\exp(v')\exp(-w))<\phi(\exp(v'))+\phi(\exp(v'')\exp(-w)). \quad (28)$$

Let $u : \mathbb{R}_+ \to \mathbb{R}$ be any strictly increasing and concave function such that $u(0) = 0$, $u(x') = \exp(v')$, $u(x'') = \exp(v'')$, and $\delta = \exp(-w)$ (which will be a number in $(0, 1)$) for some $x'' > x' > 0$. (Clearly, such a function u exists.) Then (28) gives

$$\phi(u(x'')+\delta u(0))+\phi(u(0)+\delta u(x'))<\phi(u(x')+\delta u(0))+\phi(u(0)+\delta u(x'')),$$

which is a violation of stochastic impatience. □

A.7 Ordinal Dominance

Proposition 4. *If the utility function $U : \mathbb{R}^4_+ \to \mathbb{R}$ satisfies ordinal dominance, then it can be expressed in the weakly separable form* (10)*, where f is given by (19) and G is increasing, continuous, and symmetric. Conversely, if U has the form (10), where f is increasing and continuous and G is increasing, continuous, and symmetric, then U satisfies ordinal dominance.*

Proof. Suppose U satisfies ordinal dominance. Let $f : \mathbb{R}^2_+ \to \mathbb{R}$ be defined by (19). Since f is continuous and strictly increasing, its range $\mathcal{R}(f) = \{f(a, b) : (a, b) \in \mathbb{R}^2_+\}$ coincides with $\{f(z, z) : z \in \mathbb{R}_+\}$. The function $G : \mathcal{R}(r) \to \mathbb{R}$ is defined by

$$G(f(a,a),f(b,b))=U((a,a),(b,b))$$

for all a, $b \in \mathbb{R}_+$. To verify that this definition is consistent note that f is one-to-one on the restricted domain of points of the form $(z, z) \in \mathbb{R}^2_+$. Let us verify that (10) holds. Let $((a, b),(a', b')) \in \mathbb{R}^4_+$ be an arbitrary bundle. By the fact that f is continuous and increasing there exists d and d' so that $f(a,b) = f(d,d)$ and $f(a',b') = f(d',d')$. As U satisfies ordinal dominance we have $U((a,b),(a',b')) = U((d,d),(a',b')) = U((d,d),(d',d'))$. Finally, we see from our definition of G that $U((a,b),(a',b')) = G(f(a,b),(a',b'))$ which confirms that U has a weakly separable representation.

Suppose U has a representation of the form in (10) where f is increasing and continuous and G is increasing, continuous, and symmetric. We prove that U satisfies ordinal dominance. The fact that G is symmetric ensures that U satisfies lottery equivalence. Next, suppose we have $U((a,b),(a,b)) \geq U((a',b'),(a',b'))$ for some a, a', b, b'. As f and g are increasing we obtain $f(a,b) \geq f(a',b')$. Let a'' and b'' be non-negative numbers; the fact that g and f are increasing guarantees that $U((a,b),(a'',b'')) = G(f(a,b), f(a'',b'')) \geq G(f(a',b'), f(a'',b'')) = U((a',b'),(a'',b''))$ as required. □

References

Abdellaoui, M., Bleichrodt, H., L'Haridon, O., and Paraschiv, C. Is there one unifying concept of utility? an experimental comparison of utility under risk and utility over time. *Management Science*, 59:21532169, 09 2013.

Abdellaoui, M., Kemel, E., Panin, A., and Vieider, F. M. Measuring time and risk preferences in an integrated framework. *Games and Economic Behavior*, 115:459 – 469, 2019.

Afriat, S. The construction of a utility function from expenditure data. *International Economic Review*, 8, 1967.

Afriat, S. Efficiency estimation of production functions. *International Economic Review*, 13:568–598, 1972.

Ahn, D., Choi, S., Gale, D., and Kariv, S. Estimating ambiguity aversion in a portfolio choice experiment. *Quantitative Economics*, 5:195–223, 2014.

Anderhub, V., Güth, W., Gneezy, U., and Sonsino, D. On the interaction of risk and time preferences: An experimental study. *German Economic Review*, 2(3):239–253, 2001.

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

Andersen, S., Harrison, G. W., Lau, M. I., and Rutström, E. E. Multiattribute utility theory, intertemporal utility, and correlation aversion. *International Economic Review*, 59(2):537–555, 2018.

Andreoni, J. and Sprenger, C. Risk preferences are not time preferences. *American Economic Review*, 102(7):3357–3376, 2012a.

Andreoni, J. and Sprenger, C. Estimating time preferences from convex budgets. *American Economic Review*, 102(7):3333–56, December 2012b.

Andreoni, J., Feldman, P., and Sprenger, C. A stream of prospects or a prospect of streams: On the evaluation of intertemporal risks. *NBER Working Paper*, (w24075), November 2017.

Augenblick, N., Niederle, M., and Sprenger, C. Working over Time: Dynamic Inconsistency in Real Effort Tasks *. *The Quarterly Journal of Economics*, 130(3):1067–1115, 05 2015.

Baucells, M. and Heukamp, F. Common ratio using delay. *Theory and Decision*, 68: 149–158, 02 2010.

Baucells, M. and Heukamp, F. H. Probability and time trade-off. *Management Science*, 58(4):831–842, 2012.

Bommier, A. Risk Aversion, Intertemporal Elasticity of Substitution and Correlation Aversion. *Economics Bulletin*, 4(29):1–8, 2007.

Chakraborty, A., Halevy, Y., and Saito, K. The relation between behavior under risk and over time. *American Economic Review: Insights*, 2(1):1–16, March 2020.

Cheung, S. L. Risk preferences are not time preferences: On the elicitation of time preference under conditions of risk: Comment. *American Economic Review*, 105(7): 2242–60, July 2015.

Chew, S. H. and Epstein, L. G. Nonexpected utility preferences in a temporal framework with an application to consumption-savings behaviour. *Journal of Economic Theory*, 50(1):54 – 81, 1990.

Choi, S., Fisman, R., Gale, D., and Kariv, S. Consistency and heterogeneity of individual behavior under uncertainty. *American Economic Review*, 97(5):1921–1938, 2007.

Choi, S., Kariv, S., Wieland, M., and Silverman, D. Revealing preferences graphically: An old method gets a new tool kit. *American Economic Review*, 104(6):1518–1550, 2014.

Coble, K. and Lusk, J. At the nexus of risk and time preferences: An experimental investigation. *Journal of Risk and Uncertainty*, 41:67–79, 08 2010.

Debreu, G. Cardinal utility for even-chance mixtures of pairs of sure prospects. *The Review of Economic Studies*, 26(3):174–177, 1959.

DeJarnette, P., Dillenberger, D., Gottlieb, D., and Ortoleva, P. Time lotteries and stochastic impatience. Working Paper, August 2019.

Dillenberger, D., Gottlieb, D., and Ortoleva, P. Separation of time and risk preferences. Working Paper, November 2018.

Duffie, D. and Epstein, L. G. Stochastic differential utility. *Econometrica*, 60(2):353–394, 1992.

Epper, T. and Fehr-Duda, H. Risk preferences are not time preferences: Balancing on a budget line: Comment. *American Economic Review*, 105(7):2261–71, July 2015.

Epper, T. and Fehr-Duda, H. The missing link: Unifying risk taking and time discounting. *Working Paper*, October 2018.

Epper, T., Fehr-Duda, H., and Bruhin, A. Viewing the future through a warped lens: Why uncertainty generates hyperbolic discounting. *Journal of Risk and Uncertainty*, 43(3):169–203, 2011.

Epstein, L. G. and Tanny, S. M. Increasing generalized correlation: A definition and some economic consequences. *The Canadian Journal of Economics / Revue canadienne d'Economique*, 13(1):16–34, 1980.

Epstein, L. G. and Zin, S. E. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4): 937–969, 1989.

Fagerland, M. W., Lydersen, S., and Laake, P. Recommended tests and confidence intervals for paired binomial proportions. *Statistics in Medicine*, 33(16):2850–2875, 2014.

Ferecatu, A. and Öncüler, A. Heterogeneous risk and time preferences. *Journal of Risk and Uncertainty*, 53:1–28, 2016.

Fishburn, P. *Utility Theory for Decision Making*. Wiley, New York, 1970.

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

Halevy, Y. Strotz meets allais: Diminishing impatience and the certainty effect. *American Economic Review*, 98(3):1145–62, June 2008.

Halevy, Y. Some comments on the use of monetary and primary rewards in the measurement of time preferences, April 2014. URL `https://faculty.arts.ubc.ca/yhalevy/monetary_primary.pdf`.

Halevy, Y., Persitz, D., and Zrill, L. Parametric recoverability of preferences. *Journal of Political Economy*, 126(4):1558–1593, July 2018.

Karady, G. G. The effect of temporal risk aversion on liquidity preference. *Journal of Financial Economics*, 10(4):467 – 483, 1982.

Keren, G. and Roelofsma, P. Immediacy and certainty in intertemporal choice. *Organizational Behavior and Human Decision Processes*, 63(3):287 – 297, 1995.

Kihlstrom, R. E. and Mirman, L. Risk aversion with many commodities. *Journal of Economic Theory*, 8(3):361–388, 1974.

Kreps, D. M. and Porteus, E. L. Temporal resolution of uncertainty and dynamic choice theory. *Econometrica*, 46(1):185–200, 1978.

Kubler, F., Selden, L., and Wei, X. Time consistency, temporal resolution indifference and the separation of time and risk. Working Paper, April 2019.

McClure, S. M., Laibson, D. I., Lowenstein, G., and Cohen, J. D. Separate neural systems value immediate and delayed monetary rewards. *Science*, 306:503–507, 2004.

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

Miao, B. and Zhong, S. Risk preferences are not time preferences: Separating risk and time preference: Comment. *American Economic Review*, 105(7):2272–86, July 2015.

Nishimura, H., Ok, E. A., and Quah, J. K.-H. A comprehensive approach to revealed preference theory. *American Economic Review*, 107(4):1239–63, April 2017.

Onay, S. and Öncüler, A. Intertemporal choice under timing risk: An experimental approach. *Journal of Risk and Uncertainty*, 34(2):99–121, 2007.

Polisson, M., Quah, J., and Renou, L. Revealed preferences over risk and uncertainty. Working Paper 1706, University of St Andrews, 9 2017.

Quiggin, J. A theory of anticipated utility. *Journal of Economic Behavior & Organization*, 3(4):323 – 343, 1982.

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

Richard, S. F. Multivariate risk aversion, utility independence and separable utility functions. *Management Science*, 22(1):12–21, 1975.

Selden, L. A new representation of preferences over "certain x uncertain" consumption pairs: The "ordinal certainty equivalent" hypothesis. *Econometrica*, 46(5):1045–1060, 1978.

Starmer, C. Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk. *Journal of Economic Literature*, 38(2):332–382, June 2000.

Van den Heuvel, S. Temporal risk aversion and asset prices. *SSRN Electronic Journal*, 02 2008.

Varian, H. R. The non-parametric approach to demand analysis. *Econometrica*, 50(4): 945–973, July 1982.

Varian, H. R. Goodness-of-fit in optimizing models. *Journal of Econometrics*, 46:125–140, 1990.

Wakker, P. P. Message to referees who want to embark on yet another discussion of the random-lottery incentive system for individual choice, January 2007. URL https://personal.eur.nl/wakker/miscella/debates/randomlinc.htm.

Weber, B. J. and Chapman, G. B. The combined effects of risk and time on choice: Does uncertainty eliminate the immediacy effect? does delay eliminate the certainty effect? *Organizational Behavior and Human Decision Processes*, 96(2):104 – 118, 2005.

Online Appendix: Experimental Instructions

This is an experiment on decision-making. Research foundations have provided funds for conducting this research. Your payoffs will depend partly on your decisions and partly on chance. It will not depend on the decisions of the other participants in the experiment. Please pay careful attention to the instructions as a considerable amount of money is at stake.

Your participation in the experiment and any information about your payoffs will be kept strictly confidential. In order to keep your decisions private, please do not reveal your choices to any other participant. Please do not talk with anyone during the experiment. If you have any questions, please raise your hand to ask the experimenters at any time.

The entire experiment should be completed within an hour. At the end of the experiment you will be paid privately. At this time, you will receive 10 dollars as a participation fee. The payment will be delivered to you in two times: 5 dollars 7 days later (26/6/2015) and 5 dollars 63 days later (21/08/2015). Details of how you will make decisions and receive payments will be provided below.

During the experiment we will speak in terms of experimental tokens instead of dollars. Your payoffs will be calculated in terms of tokens and then translated at the end of the experiment into dollars at the rate of **5 Tokens = 1 Dollar**.

Part I Decision Task

In the first part of the experiment, you will participate in 41 independent decision problems that share a common form. Here we describe in details the process that will be repeated in all decision tasks. In each decision task you will be asked to allocate 100 tokens among different accounts. There are four accounts Y_1, Y_2, Y_3 and Y_4.

	7 days later	63 days later
Head	$p_1 \times Y_1$	$p_2 \times Y_2$
Tail	$p_3 \times Y_3$	$p_4 \times Y_4$

You allocate 100 tokens for Y_1, Y_2, Y_3 and Y_4. <u>You can allocate any amount between 0 and 100 (including 0 and 100) to each account, and the amounts in four accounts must sum up to 100</u>. This Decision Table means that you toss a coin: if it is Head, you will receive $p_1 \times Y_1$ tokens 7 days later and $p_2 \times Y_2$ tokens 63 days later; if it is Tail, you will receive $p_3 \times Y_3$ tokens 7 days later and $p_4 \times Y_4$ tokens 63 days later. The factors p_1, p_2, p_3 and p_4 may vary for different Decision Tables.

Example. Please allocate 100 tokens in the following Decision Table.

	7 days later	63 days later
Head	1×______	2 × ______
Tail	3 × ______	4 × ______

Suppose you make the following decision.

	7 days later	63 days later
Head	1 ×___25___	2 × ___25___
Tail	3 × ___25___	4 × ___25___

It means that you will receive 25 (1 X 25) tokens 7 days later and 50 (2 X 25)tokens 63 days later if the coin toss is Head; and you will receive 75 (3 X 25) tokens 7 days later and 100 (4 X25)tokens 63 days later if the coin toss is Tail.

Suppose you make the following decision.

	7 days later	63 days later
Head	1 ×___0___	2 × ___0___
Tail	3 × ___50___	4 × ___50___

It means that you will receive 0 (1 X 0) tokens 7 days later and 0 (2 X 0)tokens 63 days later if the coin toss is Head; and you will receive 150 (3 X 50) tokens 7 days later and 200 (4 X50)tokens 63 days later if the coin toss is Tail.

Below is a screen shot that you will see for you to enter your choice.

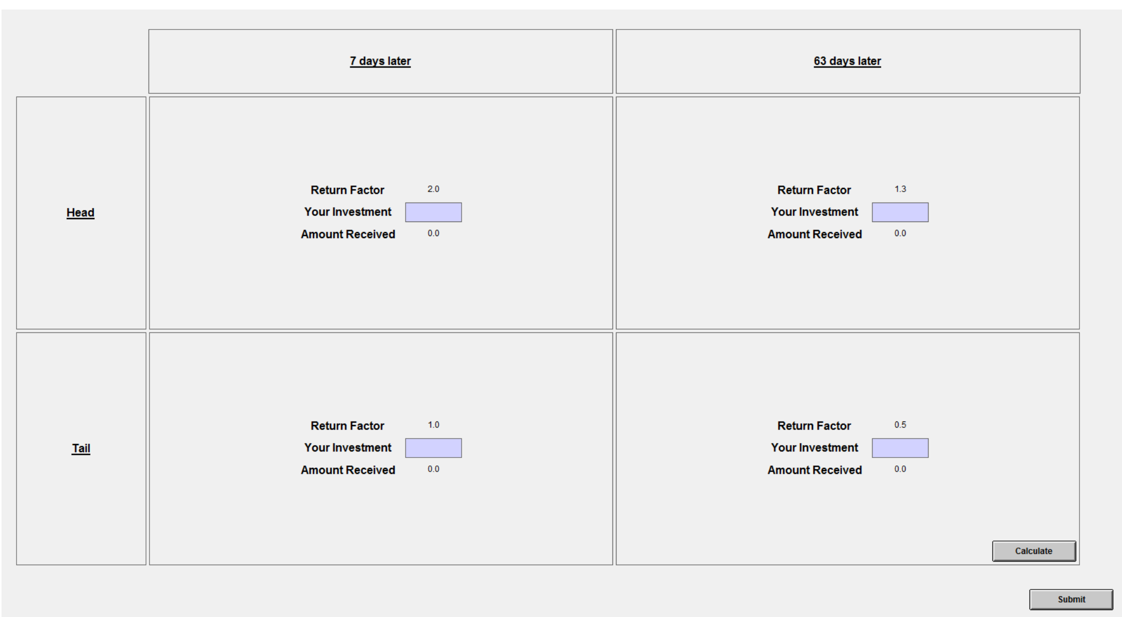

After you enter your choice, you press the “calculate” button, the computer program will help you calculate the earnings in each account as follows.

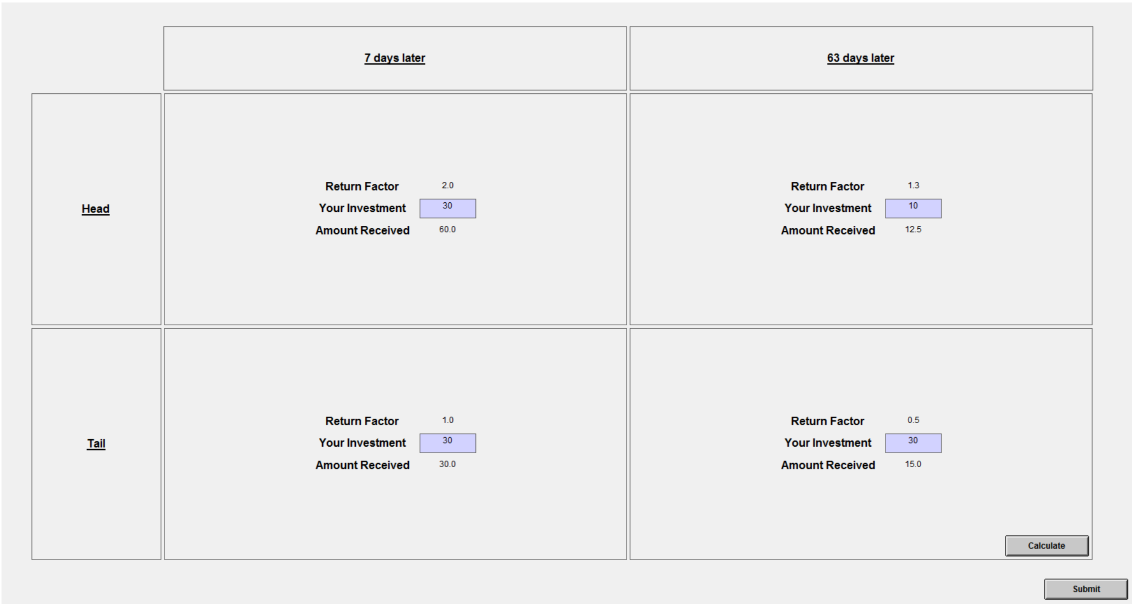

If you are sure about your decision, you can press the "submit" button to go to the next decision. If you are not sure about your decision, you can change the numbers and calculate again.

Part II Decision Task

In the second part of the experiment, you will participate in 11 independent decision problems where there are two accounts Y_1 and Y_2 as shown below.

	7 days later	63 days later
Head/Tail	$p_1 \times Y_1$	$p_2 \times Y_2$

In this case, you allocate 100 tokens for Y_1 and Y_2. This Decision Table means that: you will receive $p_1 \times Y_1$ tokens 7 days later and $p_2 \times Y_2$ tokens 63 days later regardless of the result of the coin toss. <u>In other words, you receive the tokens for sure.</u>

Similarly, below is a screen shot that you will see for you to enter your choice.

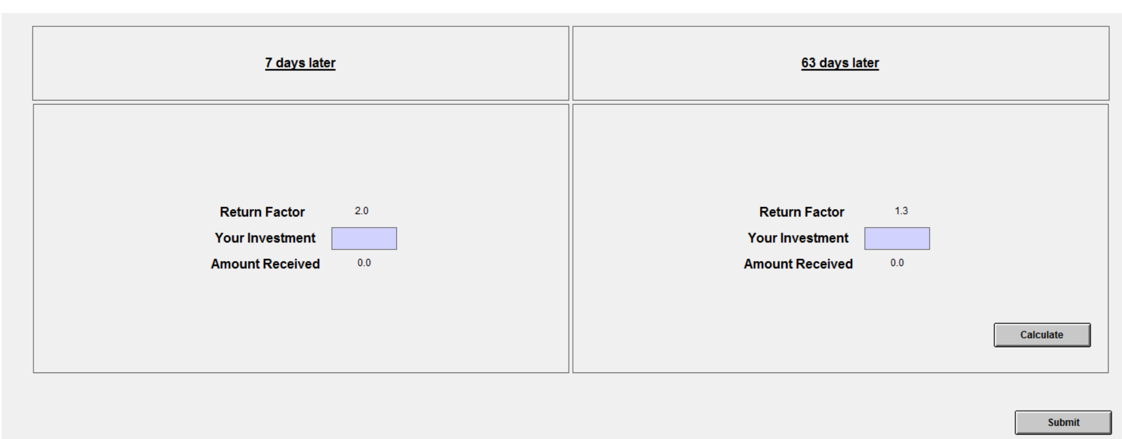

After you enter your choice, you press the "calculate" button, the computer program will help you calculate the earnings in each account as follows.

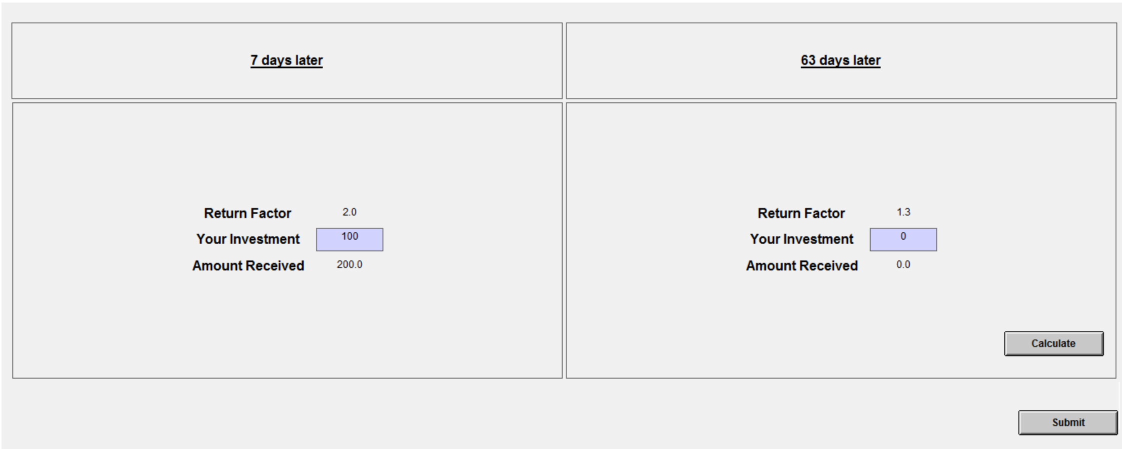

If you are sure about your decision, you can press the "submit" button to go to the next decision. If you are not sure about your decision, you can change the numbers and calculate again.

Earning

Your earnings in the experiment are determined as follows. At the end of the experiment, we will randomly select one decision from each participant to carry out (that is, 1 out of 52). The round selected depends solely upon chance. For each participant, it is equally likely that any round will be chosen. If the round selected is from Part I, your payment in two time points will be determined by a coin tossed by you and by your choice in that round. If the round selected is from Part II, your payment in two time points will be determined by your choice in that round.

At the end of the experiment, the tokens will be converted into money. Each token will be worth $0.2 (5 Tokens = $1). Your final earnings in the experiment will be your earnings realized by a coin toss in the round selected plus the $10 show-up fee ($5 in each of the two time points).

IMPORTANT: We will sign you two post-dated cheques with the specified dates (7 days later and 63 days later) at the end of today's experiment. Under Singapore banking practices, a post-dated cheque can be cashed only on or within 6 months of the specified date. It is very IMPORTANT that you do not try to cash before the date of the cheque, since you will not be able to get the money, and it will also incur a $40 loss for the experimenter.

Exercises

In order to make sure that you understand the decision making problem, we would like to complete the following exercises. After everyone finishes the exercise, we will explain the answers.

Exercise 1: Suppose you make the following decision.

	7 days later	63 days later

Head	1 ×___25___	2 × ___25___
Tail	3 × ___25___	4 × ___25___

It means that you toss a coin: if it is Head, you will receive ______ tokens 7 days later and ______ tokens 63 days later; if it is Tail, you will receive ______ tokens 7 days later and ______ tokens 63 days later.

In terms of money, if it is Head; you will receive ______+5 dollars 7days later and ______+5 dollars 63 days later; if it is Tail, you will receive ______+5 dollars 7days later and ______+5 dollars 63 days later.

Exercise 2: Suppose you make the following decision.

	7 days later	63 days later
Head	1 ×___25___	1.5 × ___30___
Tail	0.5 × ___20___	1 × ___25___

It means that you toss a coin: if it is Head, you will receive ______ tokens 7 days later and ______ tokens 63 days later; if it is Tail, you will receive ______ tokens 7 days later and ______ tokens 63 days later.

In terms of money, if it is Head; you will receive ______+5 dollars 7days later and ______+5 dollars 63 days later; if it is Tail, you will receive ______+5 dollars 7days later and ______+5 dollars 63 days later.

Exercise 3: Suppose you make the following decision.

	7 days later	63 days later
Head/Tail	1 ×___20___	0.5 × ___80___

It means that you toss a coin: if it is Head, you will receive ______ tokens 7 days later and ______ tokens 63 days later; if it is Tail, you will receive ______ tokens 7 days later and ______ tokens 63 days later.

In terms of money, if it is Head; you will receive ______+5 dollars 7days later and ______+5 dollars 63 days later; if it is Tail, you will receive ______+5 dollars 7days later and ______+5 dollars 63 days later.