Discounting Behavior: A Reconsideration

by

Steffen Andersen, Glenn W. Harrison, Morten Lau & E. Elisabet Rutström†

March 2014

ABSTRACT.

We re-evaluate the theory, experimental design and econometrics behind claims that individuals exhibit non-constant discounting behavior. Theory points to the importance of controlling for the non-linearity of the utility function of individuals, since the discount rate is defined over time-dated utility flows and not flows of money. It also points to a menagerie of functional forms to characterize different types of non-constant discounting behavior. The implied experimental design calls for individuals to undertake several tasks to allow us to identify these models, and to several treatments such as multiple horizons and the effect of allowing for a front end delay on earlier payments. The implied econometrics calls for structural estimation of the theoretical models, allowing for joint estimation of utility functions and discounting functions. Using data collected from a representative sample of 413 adult Danes in 2009, we draw surprising conclusions. Assuming an exponential discounting model we estimate discount rates to be 9% on average. We find no evidence to support quasi-hyperbolic discounting or "fixed cost" discounting, and only modest evidence to support other specifications of non-constant discounting. Furthermore, the evidence for non-constant discounting, while statistically significant, is not economically significant in terms of the size of the estimated discount rates. We undertake extensive robustness checks on these findings, including a detailed review of the previous, comparable literature.

† Department of Economics, Copenhagen Business School, Copenhagen, Denmark (Andersen and Lau); Department of Risk Management & Insurance and Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University, USA (Harrison); Durham University Business School, Durham University (Lau); and Robinson College of Business, Georgia State University, USA (Rutström). E-mail contacts: sa.eco@cbs.dk, gharrison@gsu.edu, mla.eco@cbs.ac.uk and erutstrom@gmail.com. We thank Yoram Halevy for comments, and the U.S. National Science Foundation for research support under grants NSF/HSD 0527675 and NSF/SES 0616746, the Danish Social Science Research Council for research support under project 275-08-0289, and the Carlsberg Foundation under grant 2008_01_0410. Additional appendices are available in *CEAR Working Paper 2011-03* available at http://cear.gsu.edu/papers.

Table of Contents

Different assumptions about individual discounting behavior generate significant differences in the understanding of behavior in a wide range of settings. Theorists from economics and psychology have now offered a wide range of specifications of discounting functions that match *a priori* criteria, anecdotal empirical evidence, and in some cases rigorous empirical testing. We offer a systematic and structural evaluation of most of the major alternatives.

Our approach is structural in the sense that we design experiments that allow us to jointly estimate the utility function and discounting function that individuals are assumed to use to make observed choices. We also allow for decisions to be made over shorter horizons and longer horizons, and with or without a "front end delay" on the earliest option. One of the most interesting features of the alternative specifications is the manner in which they allow short-term discounting behavior to vary, in a sense to be made clear, from longer-term behavior. Many of the earlier generation of specifications, such as the Exponential, Hyperbolic and Quasi-Hyperbolic discounting models, constrained these behaviors in ways that later specifications relax. But many of these extensions have not been evaluated in the same setting as the traditional models, nor have they been evaluated in a manner that allows several discounting models to characterize the population.

Our approach is systematic in the sense that we consider a wide range of discounting functions that characterize different aspects of the decision-making process. We do not constrain the range of discounting functions that we evaluate based on *a posteriori* inferences from other experiments or hypothetical surveys. Although this methodological approach has been productive by generating a wide range of flexible functional forms, we want to avoid it because it requires that one accepts every empirical inference that is used to characterize the discounting function. We do not believe that the behavioral landscape is as settled as some would claim, or that every such inference is well-founded in experiments that meet the usual standards of experimental economics. Our approach is to consider a range of discounting functions that span the main alternatives, and for reasons that are broadly appealing on *a priori* grounds.

In section 1 we review the alternative theoretical models that have been proposed, and settle on a list of major exemplars of the different types of models. We assume expected utility theory (EUT) for this initial purpose, and we consider alternatives as a robustness check later.[1] In section 2 we use these theoretical structures to guide the design of a series of experiments that will allow us to identify the core parameters of the latent structural models. We also discuss our specific experiments, conducted throughout Denmark in 2009 using a representative sample of the adult Danish population. In section 3 we review the econometric models used to estimate the core parameters of the models. We also explain how finite mixture models can be used to evaluate the heterogeneity of discounting behavior in the population. Section 4 contains basic results, and explores variations in some of the maintained assumptions of our basic results.

Our results are clear, and surprising. We find no support for Quasi-Hyperbolic specifications. We do find evidence in favor of flexible Hyperbolic specifications and other non-standard specifications, but with modest variations in discount rates compared to those often assumed. We find that a significant portion of the Danish population uses Exponential discounting, even if it is not the single model that best explains observed behavior.

Given the contrary nature of our findings, in terms of the received empirical wisdom, section 5 contains a systematic cataloguing of the samples, experimental procedures, and econometric procedures of the evidence for Quasi-Hyperbolic and non-constant discounting. We conclude that the evidence needed reconsideration.

One important robustness check is to see if the absent showing for the Quasi-Hyperbolic model is attributable to our population being the entire adult Danish population, rather than university students. Although it is apparent that the wider population is typically of greater interest, virtually all

[1] The logic of our approach applies to non-EUT models, since all we require is some measure of the concavity of the utility function. Those measures might be expected to be quantitatively different under EUT and non-EUT models, but our approach is quite general as a matter of theory.

prior experimental evidence that we give credence to comes from convenience samples of university students. We find that there is indeed a difference in the elicited discount rates with (Danish) university students, but that they do not exhibit statistically significant evidence of declining discount rates. The size of the discount rates for shorter time horizons is greater than that of the general population, but much smaller than the received wisdom suggests.

1. Theory

We define the discount factor for a given horizon τ to be the scalar D that equates the utility of the smaller level of income y received at time t with the larger level of income Y received at time $t+\tau$:

$$U(y_t) = D\ U(Y_{t+\tau}) \tag{0}$$

for some utility function U(.) and $y < Y$. We assume that the same utility function is used to evaluate income at time t and income at time $t+\tau$; we discuss this assumption later. This general definition of D permits the special case, much studied in the experimental literature, in which U(.) is linear. The non-linear case is of great empirical significance for inferences about discount rates, as demonstrated by Andersen, Harrison, Lau and Rutström (AHLR) [2008a]. There is nothing in (0) that restricts us to EUT, and indeed non-EUT specifications are considered later. We define utility over income and not directly over consumption flows or wealth, and discuss the implications of that specification later.

The discount factor for the Exponential (E) specification is defined as

$$D^E(t) = 1/(1+\delta)^t \tag{1}$$

for $t \geq 0$, and where the discount rate d is simply $d^E(t) = \delta$. Although these characterizations are abstract, we view the discount rate on an annualized basis throughout. The key feature of this model, of course, is that the discount rate is a constant over time. The percentage rate at which utility today and utility tomorrow is discounted is exactly the same as the rate at which utility in 7 days and utility in 8 days is discounted. The debate over climate change has reminded us all that, with this specification, even small discount rates can lead to very low weight being placed on longer-term future consequences.

The discount factor for the Quasi-Hyperbolic (QH) specification is defined as

$$D^{QH}(t) = 1 \qquad \text{if } t = 0 \tag{2a}$$
$$D^{QH}(t) = \beta/(1+\delta)^t \qquad \text{if } t > 0 \tag{2b}$$

where $\beta<1$ implies quasi-hyperbolic discounting and $\beta=1$ is exponential discounting. Although the δ in (2b) may be estimated to be a different value than the δ in (1), or other specifications below, we use the same notation to allow comparability of functional forms. The defining characteristic of the QH specification is that the discount factor has a jump discontinuity at t=0, and it is thereafter exactly the same as the E specification. The discount rate for the QH specification is the value of $d^{QH}(t)$ that solves $D^{QH}(t) = 1/(1+d^{QH})^t$, so it is $d^{QH}(t) = [\ \beta/(1+\delta)^t\ ]^{(-1/t)} - 1$ for $t>0$. Thus for $\beta<1$ we observe a sharply declining discount rate in the very short run, and then the discount rate asymptotes towards δ as the effect of the initial drop in the discount factor diminishes. The drop β can be viewed as a fixed utility cost of discounting anything relative to the present, since it does not vary with the horizon t once $t>0$. The QH specification was introduced by Phelps and Pollak [1968] for a social planning problem, and applied to model individual behavior by Elster [1979; p.71] and then Laibson [1997].

There are alternative ways to think of the fixed cost of discounting. Instead of thinking of the fixed cost as a percentage of the principal, one could think of it as a fixed monetary amount. The discount factor for the resulting Fixed Cost (FC) specification Benhabib, Bisin and Schotter [2010] is

$$D^{FC}(t) = 1 \qquad \text{if } t = 0 \tag{3a}$$
$$D^{FC}(t) = \beta\ [1 - (1-\theta)\delta t]^{(1/(1-\theta))} - (b/y_t) \qquad \text{if } t > 0 \tag{3b}$$

where $\beta<1$ indicates that there is a quasi-hyperbolic component to discounting, $b>0$ indicates that there is a fixed *monetary* cost component to discounting, and θ allows a wide range of discounting functions since $\theta=1$ (with $\beta=1$ and $b=0$) implies exponential discounting, $\theta=2$ (with $\beta=1$ and $b=0$) implies a form of hyperbolic discounting. The discount rate for the FC specification is $d^{FC}(t) = [\ \beta\ (1 - (1-\theta)\delta t)^{(1/(1-\theta))} - (b/y_t)\ ]^{(-1/t)} - 1$ for $t>0$.

There have been whole families of "hyperbolic" specifications of the discounting function. The simplest specification assumes a discount factor given by $D^{H1}(t) = 1/t$ with discount rates $d^{H1}(t) = t^{(1/t)} -$

1. The H1 specification was proposed in this manner by Ainslie [1975; p.472, Figure 3]. A slight generalization is

$$D^{H2}(t) = 1/(1+ K\ t) \tag{4}$$

for some parameter K, with discount rates $d^{H2}(t) = (1 + K\ t)^{(1/t)} - 1$, proposed by Mazur [1984; p.427]. Skipping over several additional variants documented in Appendix D (available online), a final hyperbolic generalization due to Loewenstein and Prelec [1992; p. 580] is

$$D^{H5}(t) = [1/(1+\alpha t)^{(\beta/\alpha)}] \tag{5}$$

for α, $\beta > 0$, and with discount rates $d^{H5}(t) = (1 + \alpha t)^{(\beta/\alpha t)} - 1$.

A flexible specification is based on the Weibull (W) distribution from statistics[2], and is defined as

$$D^{W}(t) = \exp(-\acute{r} t^{(1/\acute{s})}) \tag{6}$$

for $\acute{r}>0$ and $\acute{s}>0$. For $\acute{s}=1$ this collapses to the E specification, and hence the parameter $\acute{s}$ can be viewed as reflecting the "slowing down" or "speeding up" of time as perceived by the individual. This specification is due to Read [2001; p.25, equation (16)], although he noted (p.25, equation (15)) that the same point about time perception was implicit in the earlier hyperbolic generalization (5).[3] The discount rate at time t in this specification is then $d^{W}(t) = \exp(\acute{r} t^{(1-\acute{s})/\acute{s}}) - 1$.

For all of the formal specifications, there are some major themes that differentiate discounting models. For our purposes we want to focus on the exemplars of each approach, to avoid distraction with the specifics of each formulation. Obviously the E model (1) should be included as a benchmark, and the QH model (2a)-(2b) because of its popularity in behavioral economics. For the same reason, the FC model (3a)-(3b) should be considered. Within the family of "smooth" non-constant discounting models, (4) and (5) are canonical in psychology, and the W specification (6) is attractive and flexible.

[2] Any probability density function f(t) defined on [0, ∞) can form the basis of a discounting function by taking the integral of f(t) between t and ∞.

[3] The W specification is the same as the simple functional form proposed in Prelec [2004; p. 526] and applied in Ebert and Prelec [2007; p. 1424ff.] and AHLR [2008a; p. 607].

2. Experiments

There are several critical components of experimental procedures that need to be addressed when eliciting choices over time-dated monetary flows. Some are behavioral, and some are theory-driven. These components guide the specific experimental design we developed.

A. Essential Characteristics of the Experiments

The first consideration is the importance of the tradeoffs being presented in a transparent manner to subjects, rather than as a jumble of different principal amounts, horizons, front end delays, and implied interest rates. The "multiple price list" procedure for discount rate choices that was proposed by Coller and Williams [1999] is an important advance here. In this procedure the individual gets to choose between a list of options that provide a principal at some sooner date, and a larger amount of money at some future date. The list is ordered in increasing order of the larger amounts of money, to make it easy for the individual to see the tradeoffs. The intuitive aspect of this presentation is that no subject would be expected to defer payment for the first rows, where the implied return is negligible, but that every subject might be expected to defer in the last rows, where the implied return is large. Of course, "negligible" and "large" are in the eyes of the decision-maker, but annualized interest rates of less than a percentage point or more than 100 percentage points would be expected to generally fit the bill.

The second consideration, and related to the need to provide a cognitively transparent task, is the provision of annualized interest rates implied by each alternative. In many countries such rates are required to be provided as part of a regulatory requirement for most consumer loans, but one might also provide them in order to avoid testing hypotheses about whether individuals can calculate them concurrently with the effort to elicit their preferences. On the other hand, there are many settings in which real decisions with real consequences in the future do not enjoy the cognitive benefit of having implied annualized rates displayed clearly: for example, decisions to smoke, eat bad foods, engage in unsafe sex, have children, get married or divorced, and so on. Again following Coller and Williams [1999], we evaluate the provision of annualized interest rates as a treatment and study its effect on decisions.

The third component is to control for the credibility of payment. This is addressed in large part by using payment procedures that are familiar and credible, and wherever possible by adding some formal legality to the contract between experimenter and subject to pay funds in the future. Coller and Williams [1999] and Coller, Harrison and Rutström [2012] used promises to pay by a permanent faculty member that had been legally notarized; Harrison, Lau and Williams [2002] and AHLR [2008a] conducted experiments under the auspices, and actual letterhead, of a recognized government agency. One device for controlling for credibility, albeit at some cost in terms of identifying certain discounting models, is to employ a front end delay on the sooner and later payments: one argument for this procedure is to equalize the credibility of future payment for the two dated payments used to infer discount rates.[4] On the other hand, some would argue that the credibility of payment is one component of the "passion for the present" that generates non-constant discounting behavior, and that it should not be neutered by the use of a front end delay. Moreover, and critical for the present design, if the non-constancy occurs primarily within the front end delay horizon, then one might incorrectly infer constant discounting simply because the design "skipped over it." In our design we therefore want to consider as a treatment the use of a front end delay or not.[5] For the front end delay choices, both the initial and the final rewards were shifted forward by 30 days.

The fourth component is to control for the utility of time-dated monetary flows. All experimental designs prior to AHLR [2008a] assumed that utility was linear in experimental income, and

[4] Another argument is that many choices over time naturally have a front end delay. Hence the front end delay is not as artefactual a procedure as one might think.

[5] Discounting choices without a front end delay allow identification of the β-parameter in the QH specification (2b). If the "passion for the present" is shorter than the front end delay then β is simply equal to 1. One could also use several *different* front end delays to help identify the QH specification in comparison with smoothly hyperbolic specifications.

defined discount rates in terms of monetary flows instead of utility flows. This assumption had been clearly recognized earlier, such as in Keller and Strazzera [2002, p. 148] and Frederick, Loewenstein, and O'Donoghue [2002, p. 381ff.], but the quantitative importance for inferred discount rates not appreciated. A direct application of Jensen's Inequality to (0) shows that a more concave utility function must lower inferred discount rates for given choices between the two monetary options. The only issue for experimental design then is how to estimate or induce the non-linear utility function. The approach of AHLR [2008a] was to have one experimental task to identify the utility function, another task to identify the discount rate conditional on knowing the utility function, and jointly estimate the structural model defined over the parameters of the utility function and discount rate. Thus the general principle is a recursive design, combined with joint estimation of all structural parameters so that uncertainty about the parameters defining the utility function propagates in a "full information" sense into the uncertainty about the parameters defining the discount function. Intuitively, if the experimenter only has a vague notion of what U(.) is in (0), then one cannot make precise inferences about D in (0).[6]

To see the formal role of allowing for a concave utility function, assume EUT holds for choices over risky alternatives and that discounting is exponential. A subject is indifferent between two income options M_t and $M_{t+\tau}$ if and only if

$$(1/(1+\delta)^t)\ U(\omega+M_t) + (1/(1+\delta)^{t+\tau})\ U(\omega) = (1/(1+\delta)^t)\ U(\omega) + (1/(1+\delta)^{t+\tau})\ U(\omega+M_{t+\tau}) \tag{7}$$

where $U(\omega+M_t)$ is the utility of monetary outcome M_t for delivery at time t plus some measure of background consumption ω, δ is the discount rate, τ is the horizon for delivery of the later monetary outcome at time $t+\tau$, and the utility function U is separable and stationary over time. The left hand side of equation (7) is the sum of the discounted utilities of receiving the monetary outcome M_t at time t (in addition to background consumption) and receiving nothing extra at time $t+\tau$, and the right hand side is

[6] It is possible to design experimental procedures that do not require two or more experimental tasks, and embed the identification of the utility function into one task. In the case of discount rates, examples include Andreoni and Sprenger [2012a] and Laury, McInnes, and Swarthout [2012]. We discuss each in detail in Appendix D (available online).

the sum of the discounted utilities of receiving nothing over background consumption at time t and the outcome $M_{t+\tau}$ (plus background consumption) at time t+τ. Thus (7) is an indifference condition and δ is the discount rate that equalizes the present value of the *utility* of the two monetary outcomes M_t and $M_{t+\tau}$, after integration with an appropriate level of background consumption ω. This expression also makes it clear why one needs to evaluate alternative assumptions about the level of background consumption: higher ω values increase the value of the argument of the utility function, which would lead one to expect to infer more concave utility from observed risk choices, and thus lower discount rates. We consider the effect of assuming smaller values for ω, to check if that allows more "room" for discount rates to vary with the time horizon.[7]

The existing literature suggests that the front end delay and the correction for non-linear utility are the most significant treatments in terms of their quantitative impact on elicited discount rates. Coller and Williams [1999] were the first to demonstrate the effect of a front end delay; their estimates show a drop in elicited discount rates over money of just over 30 percentage points from an average 71% with no front end delay.[8] Using the same experimental and econometric methods, and with all choices having a front end delay, Harrison, Lau and Williams [2002] estimated average discount rates over money of 28.1% for the adult Danish population. AHLR [2008a] were the first to demonstrate the effect of correcting for non-linear utility; their estimates show a drop in elicited discount rates of 15.1 percentage points from a discount rate over money of 25.2%. These results would lead us to expect discount rates around 10% with a front end delay, with a significantly higher rate when there is no front end delay.

---

[7] More generally one should consider the manner in which one characterizes the degree of asset integration between net wealth outside the laboratory tasks and the prizes offered within the laboratory. We discuss this issue in section 6.

[8] The statistical significance of the front end delay is actually not clear from their results (Table 5, p.120), in part due to 22 subjects being dropped from their sample of 199 due to missing data on one variable. However, this result is readily demonstrated with their data. Appendix B (available online) contains our re-estimation of the "interval regression" statistical model they use with their complete data set.

B. The Experimental Design

Subjects are presented with two tasks. The first task identifies individual discount rates, and the second task identifies a-temporal risk attitudes. We use tasks with real monetary incentives. Observed choices from both tasks are then used to jointly estimate structural models of the discounting function defined over utility of income and background consumption. A list of parameter values for all choices is presented in Appendix A (available online).[9]

Individual Discount Rates

Individual discount rates are examined by asking subjects to make a series of choices over two certain outcomes that differ in terms of when they will be received. For example, one option can be 1000 kroner in 30 days, and another option can be 1100 kroner in 90 days. If the subject picks the earlier option we can infer that their discount rate is above 10% for 60 days, starting in 30 days, and if the subject picks the later option we can infer that their discount rate is below 10% for that horizon and start date. By varying the amount of the later option we can identify the discount rate of the individual, conditional on knowing the utility of those amounts to this individual. One can also vary the time horizon to identify the discount rate function, and of course one can vary the front end delay. This method has been widely employed in the United States (e.g., Coller and Williams [1999]), Denmark (e.g., Harrison, Lau and Williams [2002]), Canada (e.g., Eckel, Johnson and Montmarquette [2005]), and Germany (e.g., Dohmen, Falk, Huffman and Sunde [2010]).

We ask subjects to evaluate choices over several time horizons. We consider time horizons between 2 weeks and 1 year. Each subject is presented with choices over four time horizons, and those horizons are drawn at random, without replacement, from a set of thirteen possible horizons (2 weeks,

[9] The background consumption parameter ω is important, and is set exogenously: using data from the household expenditure survey at Statistics Denmark, AHLR [2008a; p.600, Appendix D] calculate per capita consumption of private nondurable goods on an average daily basis as being equal to 118 kroner in 2003. AHLR [2008a; p.602] show that estimates of discount rates are robust to variations of ω between 50 and 200 kroner. We adjust that amount for inflation to the time of our experiments, and assume ω = 130 kroner.

and 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11 and 12 months). This design will allow us to obtain a smooth characterization of the discount rate function across the sample for horizons up to one year. We also over-sampled the first three horizons, since this very short-term is clearly of great significance for the alternative specification. Hence each subject was twice as likely to get a horizon of 2 weeks, 1 month or 2 months as any of the later horizons.[10]

We also varied the time delay to the early payment option on a between-subjects basis: roughly half of the sample had no front end delay, and the other half had a 30-day front end delay. It would be possible to consider more variations in the front end delay, but we wanted to keep the treatment as simple as possible before examining the tradeoff. Similarly, we varied the provision of implied interest rates for each choice on a between-subjects basis, and independently of the front end delay treatment. We also varied the order in which the time horizon was presented to the subject: either in ascending order or descending order.

Another treatment, inspired by the intuitive notion from Benhabib, Bisin and Schotter [2010] that individuals might require a fixed monetary cost in order to delay receipt of income, is to vary the principal. The import of the "fixed cost" idea, in contrast to the notion from the QH specification that individuals require a fixed fraction of the principal to delay receipt of income, is that one should observe less "hyperbolicky" discounting as the principal gets larger and larger. Hence the non-constant discounting from a fixed monetary cost should vanish as the principal gets larger, in contrast to the QH specification. We employ two levels of the principal on a between-subjects basis (1500 and 3000 kroner), again to assess the significance of the hypothesized fixed monetary cost of delay.

These four treatments, the front end delay, information on implied interest rates, the level of the principal, and the order of presentation of the horizon, gives a 2×2×2×2 design. The subjects were assigned at random to any one particular combination of treatments, and the sample is roughly split

[10] The shorter horizons were each chosen with probability 2/16=0.125, compared to the 1/16=0.0625 probability for each of the others.

across the 16 treatments.

It is easy to see that this design allows behavior that is inconsistent with the E discounting specification. Assume that the subject is risk neutral, and switches within the "interior" of the choice options we present (see Table A2). Then a maximally non-E choice pattern would be to switch between options that offer 45% and 50% in terms of Annual Effective Rates (AER) for the two-week horizon, and between options that offer 5% and 10% in AER for the 1 month and longer horizons. This immediately implies a sharp drop in discount rates with horizon, consistent with the QH specification. A smooth drop in implied discount rates would be more consistent with the "smooth hyperbolic" specifications than the QH specification. But there is nothing in our design that biases behavior towards finding E discounting. In this respect, the most significant treatment is the front end delay. If the discounting choices always have a front end delay, as in Harrison, Lau and Williams [2002] and AHLR [2008a], one can always claim that evidence in favor of E discounting is found because all of the non-constant discounting occurs before the front end delay, as noted earlier. Addressing this concern is precisely why we have some choices that do not have any front end delay.

Risk Attitudes

Risk attitudes were evaluated by asking subjects to make a series of choices over outcomes that involve some uncertainty. To be clear, risk attitudes are elicited here simply as a convenient vehicle to estimate the non-linear utility function of the individual. The theoretical requirement, from the definition of a discount factor in (0), is for us to know the utility function over income if we are to correctly infer the discount rate the individual used. The discount rate choices described above are not defined over lotteries.

We assume that the utility function is stable over time and is perceived *ex ante* to be stable over time. Direct evidence for the former proposition is provided by Andersen, Harrison, Lau and Rutström [2008b], who examine the temporal stability of risk attitudes in the Danish population. The second proposition is a more delicate matter. Even if utility functions are stable over time, they may not be subjectively perceived to be, and that is what matters for us to assume that it is the same utility function that appears on the left-hand side and right-hand-side of (0). When there is no front end delay, this assumption is immediate for the left-hand side of (0), but not otherwise. Whether or not individuals suffer from a "projection bias" is a deep matter, demanding more research: see Ainslie [1992; p. 144-179, §6.3], Kirby and Guastello [2001] and Loewenstein, O'Donoghue and Rabin [2003].

We also assume that the same utility function that governs decisions over risky alternatives is the one that is used to evaluate time-discounted choices. This assumption has been criticized recently, and we take up those issues in section 6.

Our design poses a series of binary lottery choices. For example, lottery A might give the individual a 50-50 chance of receiving 1600 kroner or 2000 kroner to be paid today, and lottery B might have a 50-50 chance of receiving 3850 kroner or 100 kroner today. The subject picks A or B. One series of 10 choices would offer these prize sets with probabilities on the high prize in each lottery starting at 0.1, then increasing by 0.1 until the last choice is between two certain amounts of money. In fact, these illustrative parameters and design was developed by Holt and Laury [2002] to elicit risk attitudes in the United States, and has been widely employed. Their experimental procedures provided a decision sheet with all 10 choices arrayed in an ordered manner on the same sheet; we instead used the procedures of Hey and Orme [1994], and presented each choice to the subject as a "pie chart" showing prizes and probabilities. We gave subjects 40 choices, in four sets of 10 with the same prizes. The prize sets employed are as follows: [A1: 2000 and 1600; B1: 3850 and 100], [A2: 1125 and 750; B2: 2000 and 250], [A3: 1000 and 875; B3: 2000 and 75] and [A4: 2250 and 1000; B4: 4500 and 50]. The order of these four sets was random for each subject, but within each set the choices were presented in an ordered manner, with increments of the high prize probability of 0.1.

The typical findings from lottery choice experiments of this kind are that individuals are generally averse to risk, and that there is considerable heterogeneity in risk attitudes across subjects: see Harrison and Rutström [2008a] for an extensive review. Much of that heterogeneity is correlated with observable characteristics, such as age and education level (Harrison, Lau and Rutström [2007]).

C. The Experiments

Between September 28 and October 22, 2009, we conducted experiments with 413 Danes. The sample was drawn to be representative of the adult population as of January 1, 2009, using sampling procedures that are virtually identical to those documented at length in Harrison, Lau, Rutström and Sullivan [2005]. We received a random sample of the population aged between 18 and 75, inclusive, from the Danish Civil Registration Office, stratified the sample by geographic area, and sent out 1969 invitations.[11]

With a sample of 413, on average 25.8 subjects were assigned to each of the 16 treatments for the discounting tasks. We did not develop this experimental design to estimate models at the level of the individual subject or treatment condition, although obviously we will control for these factors.

Our experiments were all conducted in hotel meeting rooms around Denmark, so that travel logistics for the sample would be minimized. Various times of day were also offered to subjects, to facilitate a broad mix of attendance. The largest session had 15 subjects, but most had fewer. The procedures were standard: Appendix A (available online) documents an English translation of the instructions, and shows typical screen displays. Subjects were given written instructions, which were also read out, and then made choices in a trainer task for tiny non-monetary rewards. The trainer task was "played out," so that the full set of consequences of each choice were clear. All interactions were by computer. The order of the block of discount rate tasks and the block of risk attitudes tasks was randomized for each session. After all choices had been made the subject was asked a series of standard

[11] That recruiting sample of 1969 subjects was drawn by us from a random sample of 50,000 adult Danes obtained from the Danish Civil Registration Office, which includes information on sex, age, residential location, marital status, and whether the individual is an immigrant. At a broad level our final sample is representative of the population: the sample of 50,000 subjects had an average age of 49.8, 50.1% of them were married, and 50.7% were female; our final sample of 413 subjects had an average age of 48.7, 56.5% of them were married, and 48.2% were female.

socio-demographic questions.

There were 40 discounting choices and 40 risk attitude choices, and each subject had a 10% chance of being paid for one choice on each block. Average payments on the first block were 201 kroner (although some were for deferred receipt) and on the second block the average was 242 kroner, for a combined average of 443 kroner. The exchange rate at the time was close to 5 kroner per U.S. dollar, so earnings averaged $91 per 2 two-hour session for these two tasks. Subjects were also paid a 300 kroner or 500 kroner fixed show-up fee, plus earnings from subsequent tasks.[12]

For payments to be made in the future, the following language explained the procedures:

> You will receive the money on the date stated in your preferred option. If you receive some money today, then it is paid out at the end of the experiment. If you receive some money to be paid in the future, then it is transferred to your personal bank account on the specified date. In that case you will receive a written confirmation from Copenhagen Business School which guarantees that the money is reserved on an account at Danske Bank. You can send this document to Danske Bank in a prepaid envelope, and the bank will transfer the money to your account on the specified date.

Payments by way of bank transfer are common in Denmark, Copenhagen Business School is well-known in Denmark, and Danske Bank is the largest financial enterprise in Denmark as measured by total assets.

3. Econometrics

Our objective is to evaluate alternative discounting functions reviewed in section 1. The approach we adopt is direct estimation by maximum likelihood of a structural model of the latent choice process in which the core parameters defining risk attitudes and discounting behavior can be estimated. The approach is an extension of the full-information maximum likelihood specification used in AHLR [2008a], of course with modifications for the specification of alternative discounting functions.[13] We

---

[12] An extra show-up fee of 200 kroner was paid to 35 subjects who had received invitations stating 300 kroner, but then received a final reminder that accidentally stated 500 kroner. The additional tasks earned subjects an average of 659 kroner, so total earnings from choices made in the session averaged 1102 kroner, or roughly $221, in addition to the fixed fee of $60 or $100.

[13] AHLR [2008a] estimate a structural version of the dual-self model of "impulse control" developed by Fudenberg and Levine [2006], in which income from the experimental task is integrated into other extra-

review the inferential logic for estimating risk attitudes and discounting behavior and detailed specifications in Appendix E (available online).

4. Results

We first examine the core estimates assuming that the treatments had no effect, and then consider what conclusions change when we consider the treatments. Numerous robustness checks to the econometric specifications are also considered.

A. Initial Estimates

Table 1 reports maximum likelihood estimates of the main discounting functions. Underlying each of these sets of estimates are models of the non-linear utility function using a constant relative risk aversion (CRRA) specification, as well as the behavioral error parameters.[14] The point estimate for relative risk aversion is robustly estimated to be 0.65 with a standard error of 0.038, and a 95% confidence interval between 0.58 and 0.73. This is completely consistent with previous findings, and of course implies a concave utility function. To check the validity of the CRRA specification, we followed Harrison, Lau and Rutström [2007] and estimated the more general EP specification. We could not reject the assumption of CRRA over the domain of prizes, although there was some evidence for very slightly decreasing RRA over that domain.

The estimates in Table 1 show *robust evidence of almost-constant discounting*. There will be statistically

---

experimental income and wealth. They assume that income from the tasks that pay out on the day of the experiment is spent in one day, and income from the discount rate tasks that pay out in the future is spent in $\lambda \geq 1$ days. The structural model is estimated from choice behavior in 2003 by a sample of adult Danes. The results show that the fit of the model is maximized when $\lambda=1$, when income from both tasks are spent over the same period of time, one day. So this extended specification provides direct support for our current specification, from a sample drawn from the same population 6 years earlier.

[14] The CRRA specification we use is $U(M+\omega)^{(1-r)}/(1-r)$ for $r \neq 1$, where *r* is the CRRA coefficient. With this functional form $r = 0$ denotes risk neutral behavior, $r > 0$ denotes risk aversion, and $r < 0$ denotes risk seeking behavior. We use a Fechner stochastic error term in our statistical models, instead of the Luce specification that we used in AHLR [2008a], for greater numerical stability if $r \approx 1$.

significant evidence of non-constant discounting in some specifications, but nothing that is as dramatic in terms of economic significance as the conventional wisdom might suggest. In other specifications, there might be evidence of non-constant discounting in point estimates, but not when one allows for the statistical uncertainty of the estimates. It is not appropriate, of course, to draw inferences from point estimates without considering their statistical precision.

The Exponential discounting model indicates a discount rate of 8.9%, where all discount rates will be presented on an annualized basis. The 95% confidence interval for this estimate is between 7.4% and 10.4%, so this is slightly lower than the 10.1% reported by AHLR [2008a] for the same population in 2003. For comparison, the Exponential discounting model assuming a linear utility function implies an 18.3% discount rate, with a 95% confidence interval between 15.5% and 21.2%, so this is lower than the estimate reported in AHLR [2008a] (25.2%, with a 95% confidence interval between 22.8% and 27.6%). We again conclude that correcting for the non-linearity of the utility function makes a significant quantitative difference to estimated discount rates.

The most striking finding from Table 1, for us, is that *there is no Quasi-Hyperbolic discounting*. The key parameter, β, is not statistically or economically significantly different from 1, and the parameter δ is close to the estimate of δ from the Exponential discounting model. The *p*-value on a test of the hypothesis that $\beta=1$ has value 0.55, although the 95% confidence interval for β is enough to see that it is not significantly different from 1.

We also see from panel C of Table 1 that the rejection of the QH specification is not due to there being a different kind of fixed cost to discounting. We reject the hypothesis from the Fixed Cost discounting model (3a) and (3b) that $\beta<1$, as one might expect from panel B, but we also find no evidence that $b>0$. Furthermore, we cannot reject the joint hypothesis that $\beta=1$ and $b=0$, with a *p*-value of 0.41.[15] Because $\theta>1$ there is some evidence for hyperbolic discounting, but the statistical significance

[15] We find essentially the same results if we estimate solely on the choices made with no front end delay. The joint hypothesis that $\beta=1$ and $b=0$ is then rejected with a *p*-value of 0.92.

is very slight. Assuming β=1 and b=0, we estimate θ to be 5.88 with a standard error of 5.33, and one cannot reject the hypothesis with such a standard error that θ=1 (*p*-value of 0.36). In effect, with β=1 and b=0 this model has collapsed to a Simple Hyperbolic model, and one may as well then estimate it and Generalized Hyperbolic models.

Panels D and E do just that. The coefficient estimates by themselves are somewhat cryptic, except for those trained in the dark art of interpreting such specifications. But the Simple Hyperbolic discounting model translates into discount rates that are 9.25% for a 1 day horizon, and only decline to 8.85% for a one year horizon; in each case the 95% confidence interval for the discount rate is roughly between 7% and 11%, so there is no evidence of significantly declining discount rates. The Generalized Hyperbolic discounting model does not improve significantly on the fit of the Simple Hyperbolic model, with similar log-likelihoods.

The Weibull discounting model in panel F allows a very different pattern of non-constant discounting, but again collapses to the Exponential model. The 95% confidence intervals on all of the implied discount rate horizons is at least between 5% and 15%, and one cannot formally reject the Exponential discounting model hypothesis that ś=1 (*p*-value of 0.73).

B. Controlling for the Treatments

To what extent is the success of the Exponential discounting model due to the front end delay, the provision of information on implied interest rates, and other procedural conditions of the experiment? Table 2 reports estimates from the Exponential and QH discounting models, allowing for binary dummy covariates to reflect the effects of these treatments on β and δ. Variable FED indicates if a 30-day front end delay was employed for the "sooner" option; INFO indicates if information on implied interest rates was provided; H_ORDER indicates if the subject was presented the horizons in increasing order (rather than decreasing order); P_HIGH indicates if the higher principal of 3000 kroner was used (rather than 1500 kroner); RA_FIRST indicates if the risk aversion task was presented before the discounting task; and FEE_HIGH indicates if the higher show-up fee of 500 kroner was used to recruit the subject (rather than 300 kroner). We note in passing that the last two treatments had no statistically or economically significant effect on elicited risk attitudes.

Focusing on the Exponential discounting model, we see that INFO and H_ORDER have a statistically significant effect on the elicited discount rate. The size of the effect in each case is large in relation to our baseline estimates of discount rates, but is not large in relation to the discount rates often reported in the literature. Providing information on implied interest rates leads to a *decrease* in the elicited discount rate of 3.6%, and using increasing horizons also leads to a decrease of 3.7%.[16] The front end delay does not affect elicited discount rates in any significant manner: although the estimated effect is positive and small (2.7%), the 95% confidence interval spans zero.[17]

Turning to the QH model, we observe statistically significant effects from the FED and H_ORDER treatments on the estimated δ. The effect of the front end delay implies an increase of the discount rate of only 3.5% if we momentarily assume β=1 to interpret the effect on δ directly as the effect on the discount rate, and the effect of increasing horizons on δ is -2.7%. The only treatment to have an effect on β is whether the risk aversion task was held first: if it was, and the discounting task came second, β is estimated to be 0.023 lower.[18]

These results suggest that our main conclusion thus far, the lack of support for the QH specification in favor of the Exponential model, appears to be robust to controls for the prime suspects

[16] These results suggest that subjects require a higher premium to delay outcomes when the time horizons are presented in descending order instead of ascending order. The monetary reward of delaying an outcome is higher for longer time horizons, and it is possible that subjects are more focused on monetary rewards of delaying outcomes than implied interest rates when they first are presented with longer horizons instead of shorter horizons, or simply that the monetary threshold at which they are willing to save is smaller when the shortest time horizon is presented first.

[17] We also estimate the Exponential model with an interaction term between the two principal amounts (P_HIGH) and the time horizons between the sooner and later payments measured on a cardinal scale. There are no significant marginal effects of the control variables on elicited discount rates, which indicates that individual discount rates are constant over time and across different principal amounts.

[18] We have also estimated the model solely on choices with no front end delay and find that the estimated coefficient on β is equal to 0.994 with a standard error of 0.016.

in terms of our elicitation procedures. Essentially the same is true for the other specifications.

One of our "treatments," in a sense, is the elicitation of discount rates over horizons extending from 2 weeks up to one year. To what extent is the lack of evidence for hyperbolicky discounting due to constant-discounting responses to longer horizons swamping non-constant responses to shorter horizons? One could simply re-weight the data to focus more on the shorter horizons, but a simpler method is simply to estimate the models with shorter horizons. Focusing on the two shortest horizons of 2 weeks and 1 month, which is roughly 25% of the data due to our deliberate over-sampling design, we do not see any deviations from constant discounting. There is no statistically significant effect on either parameter of the QH specification, and we reach the same conclusion with the Weibull discounting model.

C. Robustness Checks

Appendix F (available online) considers several robustness checks on our results.

The first is to consider a non-EUT specification of behavior with respect to the risky lotteries, and see if that changes inferences about the curvature of the utility function and hence discount rates. We model lottery choices behavior using a Rank-Dependent Utility (RDU) model, since all choices were in the gain frame, and find evidence of probability weighting. The probability weighting function is S-shaped with underweighting of small probabilities and overweighting of high probabilities. Despite the evidence of probability weighting, the vast bulk of aversion to risk derives from aversion to variability of outcomes and the utility function is more concave than under EUT. We do not find any evidence of non-constant discounting when we allow for probability weighting in the statistical models.

The second robustness check is to consider mixture models in which the observed choices over time-dated outcomes could be generated by two discounting models rather than one, with some fraction of observed choices accounted for by one model and the remaining choices accounted for by the other model. The mixture model specification jointly estimates the structural parameters of each model as well as the mixing probability between the two of them. We first consider a mixture between the Exponential and QH discounting models. The estimates attach substantial weight to the QH model, and the probability for that model is estimated to be 0.72 Thus it would appear that there is considerably more support for the QH specification, until one examines the estimated parameter values for each model. In the case of the Exponential model the discount rate is 15.3%, and for the QH model it is effectively a second Exponential specification with a discount rate of 3.3% because the estimate of β is essentially 1. Thus we see a bimodal distribution in the sample, with just over a quarter of the choices being characterized by a discount rate of 15.3% and the rest by a discount rate of 3.3%.[19]

We also consider a mixture between the Exponential and Weibull discounting models, and come to the same conclusion. In this case the sample divides into one mode with 29% of choices at a discount rate of 14.7%, and the other mode with Weibull discounting that is not statistically significantly different from Exponential (p-value = 0.21). The discount rates implied by the Weibull parameter estimates range from 7.6% for a 1 day horizon, to 5.2% for a one week horizon, to 3.9% for a one month horizon, down to 2.5% for a one year horizon. But it is important to recognize the relative statistical imprecision of the implied discount rates in this Weibull specification for the shorter horizons.

The third robustness check is to allow for observed and unobserved individual heterogeneity in behavior, both in terms of the choices over risk, and inferences about the utility function, and the choices over time-dated outcomes. We first consider observed heterogeneity by means of a set of standard socio-demographic characteristics of individuals affecting each of the structural parameters we estimate. The upper panel in Figure 1 shows the implied distribution of predicted Exponential discount

[19] AHLR [2008a] considered a mixture model between Exponential and Simple Hyperbolic discounting models and found that 72% of the observations could be characterized by the Exponential specification with a discount rate of 6.8%, and the remaining 28% of the observations could be characterized by the Simple Hyperbolic specification with K=0.33. Although the Simple Hyperbolic specification does not directly compare to the Quasi-Hyperbolic or Weibull discounting models, we find similar results in the sense that 72.4% of the observations can be characterized by a low discount rate of 3.3%, and the remaining 27.6% of the observations can be characterized by another Exponential specification with a high discount rate of 15.3%.

rates from the subjects in our sample: the mean estimate is 11.3% with a standard deviation of 3.0%. The middle panel displays the estimated population distribution of discount rates using random coefficients for unobserved heterogeneity, with a higher mean estimate of 12.1% and a larger standard deviation of 8.5%. Finally, we consider estimation at the level of the individual: the mean is 11.7 and the standard deviation is 10.2%. The distribution of estimated discount rates in the random coefficients model reflects the variation in discount rates at the individual level, which illustrates the complementarity between the two estimation methods. There is considerable variation in estimated discount rates for each of three different estimation methods, with more evidence of variation in the models that allow for unobserved heterogeneity.[20]

The estimated distributions of the β parameter in the QH model are displayed in Figure 2. We observe, in the upper panel, that the estimated β coefficient has a mean of 0.99 with a standard deviation of 0.02 when we control for observed heterogeneity. Allowing for unobserved heterogeneity with random coefficients *or* undertaking individual estimation generates a distribution with a mean of 0.99 and standard deviation of 0.03. Finally, Figure 3 shows the estimated distributions of δ in the QH model for different econometric specifications of individual heterogeneity.[21] The results are qualitatively the same as for the Exponential model in Figure 1: the mean (standard deviation) of the distribution is 10% (3%), 13% (10%), and 9% (8%) for the top, middle and bottom panels. None of the three approaches to modeling individual heterogeneity changes our basic conclusions about the best discounting model to characterize these data.

The fourth robustness check is to see if the absence of evidence for hyperbolic discounting is due to the theoretically motivated use of a concave utility function when inferring discount rates from observed choices over time-dated amounts of money. Since a concave utility function significantly

[20] The models with observable heterogeneity and random coefficients are based on responses from 308 subjects. We solved the model at the individual level for 228 of these subjects.

[21] The models with observable heterogeneity and random coefficients are based on responses from 198 subjects in the no-FED treatment. We solved the model at the individual level for 80 of those subjects.

lowers the level of discounting inferred, perhaps that means that there is simply less "headroom" for discount rates to be hyperbolic. This is easy to check by simply assuming a linear utility function for each of our specifications. We certainly infer higher discount rates, but in no case do we observe any statistically significant decline in discount rates with horizon.

Finally, we consider the effect of assuming smaller values for ω, to check if that allows more "room" for discount rates to vary with the time horizon. We check this by setting $\omega = 0$ and find that our results are robust to this variation in background consumption. Since earnings were realized immediately after each decision task, one could also integrate earnings from the first decision task with income in the second decision task. It is not immediately clear to what extent subjects would integrate income at different dates in the discounting task with earnings from a previous risk aversion task, and vice versa. This is an avenue for future research, and one option is to consider partial asset integration models, in which subjects behave as if some fraction of personal wealth or income is combined with experimental prizes in the utility function: this combination implies less than perfect substitution (Andersen, Cox, Harrison, Lau, Rutström and Sadiraj [2011]).

5. Connection to Previous Literature

Our results were a surprise to us, and the robustness checks reported above did not lead us to qualify that reaction. We fully expected to see much more "hyperbolicky" behavior when we removed the front end delay, and particularly when that was interacted with not providing the implied interest rates of each choice. We were not wedded to one hyperbolic specification or the other, and did not expect the exponential model to be completely overwhelmed by the alternatives, but we did expect to see much more non-constant discounting. We therefore re-examined the literature, and tried to draw some inferences about what might explain the apparent differences in results.

A. Reconsideration of Previous Literature

We undertook a re-examination of the previous, comparable literature that has led to the conventional wisdom of significant non-constant discounting. We ignored all hypothetical survey studies, on the grounds that the evidence is overwhelming that there *can* be huge and systematic hypothetical biases. It is simply inefficient to take the evidence from hypothetical survey studies seriously.[22] We also focused on experiments, rather than econometric inferences from naturally occurring data, because those data are easier to interpret and have generated the conventional wisdom.[23] We excluded studies that did not lend themselves to inferring a discount function.[24] Finally, we excluded any study that used procedures that were not incentive-compatible or that involved deception.[25]

Table 3 summarizes the studies we examined, and Appendix D (available online) contains more details on several of the more important studies. Our objective is not to dismiss or "discount" all evidence for non-constant discounting, but just to weigh it carefully to see if it is as monolithic as has been claimed. One conclusion that we draw is that most evidence of non-constant discounting comes from studies undertaken with students. We therefore conducted a conventional laboratory experiment, described below, using the same procedures as in our (artefactual) field experiment but with students recruited in Copenhagen.

---

[22] Harrison [2006] and Harrison and Rutström [2008b] provide surveys of the literature. We use the literature reviews of Coller and Williams [1999] and Frederick, Loewenstein and O'Donoghue [2002] as an initial guide; it should be noted that the latter list Holden, Shiferaw and Wik [1998] as using real incentives, although they did not (see p. 110).

[23] For example, Harrison [2005; §4.2] discusses at length the difficulties making robust inferences from the natural experiment studied by Warner and Pleeter [2001]. Appendix D (available online) reviews the results from one additional study of interest using naturally-occurring data.

[24] For example, experiment 3 of Read, Frederick, Orsel and Rahman [2005] was designed to test if one obtained the same results when the later horizon was presented as a real date or as a time delay. Although one might infer discounting functions from their data, the design does not lend itself to that type of inference.

[25] For example, Experiment 1 of Kirby and Maraković [1995] had both problems. They used a first-price sealed-offer auction between 3 subjects to elicit the present value of a future amount, and acknowledge that an optimal (risk-neutral) bid would be above the true valuation (just as an optimal bid for a risk-neutral agent in a first-price sealed-bid auction is below true valuation). They also conducted auctions with only 3 bidders, which makes the optimal overstatement more severe than if the auction were for many more bidders: as the number of bidders increase the mis-statement decreases quite rapidly. Furthermore, they deceived subjects and actually had them bid against simulated opponents.

Two additional conclusions from the review of the literature are the potential roles of "small stakes" and open-ended, "fill in the blank" elicitation procedures. Evaluating these characteristics of the previous literature is beyond the scope of our study, since they raise a host of behavioral and experimental issues.[26]

B. Experiments with Students

In order to determine if the evidence for non-constant discounting derives from the general focus on students samples, we replicated our field experiments with a student sample in Copenhagen recruited using standard methods.[27] The experimental tasks were identical, to ensure comparability.

Table 4 lists estimates from the student responses of the basic models in Table 1. The risk attitudes of this sample were close to those of the adult Danish population.[28] The results are intriguing, but do not change our basic conclusions. We do observe a slightly higher discount rate with the Exponential model: 10.4% compared to 8.9% for the adult population in general. And we obtain point estimates that suggest modest QH discounting (β=0.986), although the 95% confidence interval spans β=1 and we can reject the assumption of β=1 at the p=0.18 level for a two-sided test, and hence at the p=0.09 level for the appropriate one-sided test. So this is suggestive of QH discounting, but not clear

[26] The use of small stakes generates profound confounds when subjects "round" responses up to the nearest major currency unit, such as a dollar in the United States. Andersen, Harrison, Lau and Rutström [2013] show that simple rounding can explain both the "magnitude effect" and "hyperbolicky" behavior with small stakes of the scale found in most laboratory experiments with real rewards. The use of open-ended elicitation procedures, in which present values or future values are directly elicited, opens up problems of subjects comprehending the incentive compatibility of those tasks (documented in closely related settings by Rutström [1998] and Harstad [2000]). In addition, popular open-ended elicitation procedures are known to generate extremely weak incentives for subjects to respond truthfully or precisely (Harrison [1992]).

[27] Sessions were announced at two large lectures at Economics classes at the University of Copenhagen for 1st and 2nd year students. In addition, posters were put up at most of the major student dormitories. The students then had to send an email to get listed for one of the sessions. We recruited some 11 to 12 subjects for each session, and easily filled the available sessions. Our sample of 88 subjects consists of a broad array of types of students, not just Economics students.

[28] The coefficient r in the CRRA utility function is estimated to be 0.53, with a 95% confidence interval between 0.42 and 0.64. There is no evidence of varying relative risk aversion over this domain: the coefficient α in the Expo-Power utility function has a p-value of 0.19 for the hypothesis that α=0.

evidence. We find no evidence of Fixed-Cost discounting, and no evidence of Simple Hyperbolic discounting. We do observe some non-constancy of discount rates with the Weibull discounting specification, although the overall effect of the student sample is not statistically significant, as shown by the *p*-value of 0.17 on the null hypothesis that the specification is actually Exponential.[29]

C. The Probability of Discounting

Our literature review deliberately ignored studies that do not consider decisions with real monetary consequences, but there is one treatment that we do want to recognize even if it has only been addressed in studies using hypothetical survey questions: the effect of the discounting tasks being rewarded probabilistically.[30] In our experiments each subject had a 10% chance that one of their discounting tasks would be rewarded.[31]

Keren and Roelofsma [1995] demonstrated, with hypothetical tasks, a behavioral effect of this treatment on discounting behavior, specifically a reduction in the extent of hyperbolicky behavior for shorter horizons. Their first experiment illustrates their findings. Subjects were offered 100 Dutch Guilders now or 110 in 4 weeks. When the payment was not probabilistic, 82% of 60 subjects chose the

[29] The Weibull estimates in Table 4 imply discount rates of 25.7% for a 1 week horizon, with a 95% confidence interval between 4.9% and 46.6% (the estimated rates for shorter horizons are higher, but even less precisely estimated). After 2 weeks the estimated rate is 19.9% (8.0% ↔ 31.8%), after 1 month it is 14.9% (9.2% ↔ 20.6%), after 3 months it is 9.9% (7.2% ↔ 12.5%), and after 1 year it is 5.9% (2.4% ↔ 9.3%). So the discount rates for a 1 week horizon are significantly higher than those for horizons of 1 month or more, and these *pairwise* differences are quantitatively significant. For the adult population, the rates for a 1 week horizon were 9.7% (2.6% ↔ 16.8%), for a 2 week horizon they were 9.4% (4.0% ↔ 14.7%), for a 1 month horizon they were 9.0% (5.5% ↔ 12.6%), for a 3 month horizon they were 8.6% (7.0% ↔ 10.1%), and for a one year horizon they were 8.1% (5.6% ↔ 10.5%).

[30] In the economics literature, Halevy [2008] emphasizes this effect, but does not present new experimental evidence for (or against) it. Epper, Fehr-Duda and Bruhin [2010] also conducted discounting experiments with every subject being certain of one of their choices being rewarded, since their core hypotheses have to do with the effect of uncertain payoffs in conjunction with sub-additive probability weighting. They did not conduct a control experiment with some probability of the subject being paid that would allow this treatment to be studied, nor was it needed for their design purposes.

[31] Since we only paid each subject with a 10% probability in the risk aversion and discounting tasks, one could argue that the subjects made binary choices over compound lotteries in the risk aversion task, and binary choices over uncertain amounts at different dates in the discounting task. Our results are virtually the same if we take the 10% probability of being paid in the two decision tasks into account in our statistical analysis.

sooner option. But when the payment would occur with a probability of 0.9, only 54% of 70 subjects chose the sooner option, and that declined to only 39% of 100 subjects with a probability of 0.5. With a front end delay of 26 weeks, the same subjects chose the sooner option in 37%, 25% and 33% of the choices, respectively, suggesting that reducing the probability of payment to 0.5 generated results consistent with having a front end delay.

Weber and Chapman [2005] were unable to replicate these findings.[32] They used 446 students in an introductory psychology class in a between-subjects replication of the experiment described above, but with U.S. Dollar amounts instead of Dutch Guilders. With no front end delay the fraction choosing the sooner option was 61% (of 113) with a probability of 1 of payment, and 70% (of 111) with a probability of 0.5. Adding a front end delay generated comparable choices of 46% (of 109) and 51% (of 113), implying no significant effects of having probabilistic payments. Although one hesitates to pursue design differences with non-salient tasks, it is worth noting that the Keren and Roelefsma [1995] subjects were compensated for attending the session.

We consider the effect of probabilistic discounting by undertaking experiments in which we vary the exogenous probability of payment. Specifically, we conduct experiments with 28 subjects from the greater Copenhagen area in which we vary the probability of payment for the discounting task from 10% to 100%, and see if there is a difference in behavior. Of course, increasing the probability means that we need to account for the scale effects on expected rewards. In Keren and Roelefsma [1995] the stakes were kept the same, so there may be a confound of a scale effect. For example, their subjects might have been close to risk-neutral for lower stakes (hence implying higher discount rates when the

---

[32] Halevy [2008; p. 1148] notes that "Weber and Chapman [2005] replicated Keren and Roelofsma's [1995] findings." This is not completely correct. Experiment 1 of Weber and Chapman [2005] was their only direct replication of the *design* reported by Keren and Roelofsma [1995], and reproduced by Halevy [2008; p. 1148], and decisively failed to replicate the original findings. In a footnote, Halevy [2008; p. 1148, fn. 10] adds that "The reader is referred to Experiment 2 (summarized in Tables 5 and 6) in their study." But this Experiment 2 had some significant and problematic differences in design from the original, involving the use of dubious indifference-point elicitation procedures. So although it changed the *design* from the original, it did replicate the *finding* from the original.

stakes were to paid with some probability less than 1) and risk averse for higher stakes (hence implying lower discount rates when the probability of payment was closer to 1). This pattern of risk aversion is found in many laboratory settings: for example, see Holt and Laury [2002][2005] and Harrison, Johnson, McInnes and Rutström [2005]. We therefore maintain the stakes at their original levels, despite the cost of the experiments, and allow for varying risk aversion with stakes.

The change in instructions in the IDR task was simple. The original text was:

> You will have a 1-in-10 chance of being paid for one of these decisions. The selection is made with a 10-sided die. If the roll of the die gives the number 1 you will be paid for one of the 40 decisions, but if the roll gives any other number you will not be paid. If you are paid for one of these 40 decisions, then we will further select one of these decisions by rolling a 4-sided and a 10-sided die.

The new text was simply this:

> You will be paid for one of the 40 decisions. We will select one of these 40 decisions by rolling a 4-sided and a 10-sided die.

The experiments with this 100% treatment were conducted in September 2010, and used the lower principal in our baseline experiments. All other conditions were the same.

Reviewing the set of discounting models in Table 1, we find very little effect from this treatment. We first re-estimate each model with a dummy added to capture the effect of the new experiments for each discounting parameter, and we then estimate using only the new sample.

For example, for the Exponential model we first estimate δ_0 and δ_1 in $\delta = \delta_0 + \delta_1 \times C$, where C is a binary indicator variable for the 100% certain responses. There is no statistically significant effect on the discounting parameter(s) for the Exponential, Simple Hyperbolic, Fixed Cost Hyperbolic, Generalized Hyperbolic, and Weibull models.

For the QH discounting model there is an effect on the all-important β parameter.[33] It is 0.025 lower with the 100% payment treatment, and this effect has a *p*-value of 0.074. Estimating the QH

[33] The Fixed Cost Hyperbolic shows the same effect when constrained to the QH, of course, but not when it is estimated in unconstrained form.

model with the new sample we estimate β to be 0.983 with a standard error of 0.013, and δ to be 0.041 with a standard error of 0.024. The hypothesis that β=1 has a *p*-value of 0.185, so this is not statistically significant evidence in favor of the QH model.

We conclude that the effect of probabilistic discounting is non-existent or negligible in our sample, and for the specifications considered here.

6. Open Issues

Reliably inferring risk and time preferences is not easy. Despite the progress of recent years, we believe that there are several important open issues.

First, we need to allow for alternative models of decision-making over risk for some decision-makers in some settings. The identification of non-standard models of risk preferences, such as RDU or Cumulative Prospect Theory, demands careful attention to the tasks given to subjects, and is not something that we believe can be safely "folded in" with some other task, as proposed by Andreoni and Sprenger [2012a]. In a similar vein, interacting risk and time preferences, say by offering subjects a choice over time-dated lotteries, may raise deep confounds if one insists on standard, additive intertemporal utility functions. The claim that "risk preferences are not time preferences" of Andreoni and Sprenger [2012b] can be viewed as an illustration of that confound at work (see Harrison, Lau and Rutström [2013]).

Second, the implications for inferred risk attitudes of worrying about asset integration are more subtle than recent controversies over calibration might suggest. Proper identification of the extent to which subjects in experiments integrate the prizes in tasks with "outside wealth" demands unique data, and does not always lead to the conclusion that subjects have to be risk neutral over small stakes (e.g., Andersen, Cox, Harrison, Lau, Rutström and Sadiraj [2011]).

Third, the discount rates that are implied over monetary prizes and over consumption flows can be quite different, as stressed by Cubitt and Read [2007]. We do not take the view that the only possible or interesting argument of a utility function is a consumption flow, but we are certainly interested in such flows. There are methods for allowing explicitly for the relationship between monetary prizes and consumption flows, as we illustrate at length in AHLR [2008a]. There is a need for comparable experiments examining risk and time preferences over real flows, although "sips of juice" and equally contrived examples of real effort are not what we find convincing.

Related to this point, one intriguing hypothesis behind our finding that Danes tend to discount in an Exponential manner could be that our experiment served as an artefactual commitment device. The idea is that our highly structured experiments, and the formal manner in which we explained the credibility of the our making payments on specified future dates, might appeal to individuals that have difficulties making consistent intertemporal plans in terms of other choices they make. This is a hypothesis about the external validity of our findings, and the extent to which they might transfer to other, less-structured intertemporal choices. We find this hypothesis attractive *a priori*, and just note that such control and structure may be the price one pays for internal validity in claiming to have measured time preferences. Building a bridge between that internal validity and a wider class of field choices is an important challenge for future experiments.

Fourth, our approach to identification of discount rates defined over utility has always made one assumption we find problematic: that the a-temporal utility function the subject exhibits today is the same a-temporal utility function the same subject applies to evaluate future monetary prizes or consumption flows. In behavioral terms, we assume away any "projection bias," as noted earlier, and should instead use the subjectively expected utility for the future self. We do not know yet how to reliably identify the latter concept.

Finally, the level of stakes is a deeper issue than many believe. With small stakes, it is easy to demonstrate that hyperbolicky discounting and the magnitude effect can arise if subjects simply round the monetary amounts to some natural unit. AHLR [2013; p.682ff.] demonstrate this conclusion using the design of Benhabib, Bisin and Schotter [2010]. One solution, which we adopt and checked, is to use larger stakes and then have some percentage of the subject actually being paid. However, this solution might make field arbitrage opportunities more salient. In that case one should consider the effect of censoring of responses in line with borrowing and savings interest rates that the subject has available outside the experiment, as in Coller and Williams [1999] and Harrison, Lau and Williams [2002].

7. Conclusions

We do not see significantly hyperbolicky discounting behavior in adult Danes making choices of deferred monetary payments. If there is any statistically significant evidence for non-constant discounting, and there is in a fraction of the population, it entails discount rates that for many practical purposes are virtually constant.

How do we reconcile this striking finding with the received wisdom? We see nothing in our experimental procedures which might bias behavior, and that deviates in any novel manner from the types of procedures used in the past. We avoid eliciting present values in an open-ended manner, because we are suspicious of the behavioral accuracy of those responses.[34] We test for the effect of providing information on the implied interest rates we offered. We use displays of the tasks that make them relatively transparent in terms of the choice alternatives, rather than rely entirely on the ability of subjects to read numbers and words. And, obviously, we pay the subjects in a salient manner.

Our basic econometric procedures are familiar from the binary choice literature, and have a long tradition in experimental economics (e.g., Camerer and Ho [1994] and Hey and Orme [1994]). Our application of them uses parametric methods, but we are clearly flexible in terms of the discounting functions we examine. The notion of joint estimation of utility functions and discounting functions is driven by theory, and implies nothing fundamental from an econometric perspective. The application of

[34] We strongly encourage systematic studies of the effects of using discrete choice and open-ended "matching" procedures, along the lines of Ahlbrecht and Weber [1997] and Read and Roelofsma [2003], but for discounting tasks in which subjects are making salient, non-hypothetical choices.

mixture specifications to explore the robustness of our basic results is, similarly, not fundamentally novel in terms of method.

With some exceptions, noted in our literature review, the evidence of hyperbolicky behavior that meets certain minimal standards of salience and design occurs in samples of college-age students. We do not dismiss student samples as irrelevant, or the exceptions as flawed studies: our point is just that it is difficult to make inferences about behavior in general from a small student population. We provide some evidence of hyperbolicky behavior in a sample of college-age students in Denmark, but the results are not statistically significant and the quantitative extent is relatively modest in relation to the literature, even if we do view it as nonetheless economically important.

Theorists use illustrative examples of hyperbolicky behavior towards things like the "eating of potato chips" as metaphor. If it is a poor metaphor when applied to monetary choices of adult Danes over horizons of weeks and months, that means that there is an important empirical bridge to be built. What are the tasks, domains, and samples for which hyperbolic behavior might be expected to apply for significant sub-samples? The metaphor may have been stretched too far, but it refers to impulsive choices over foods and alcohol, drugs, sexual habits, driving behavior, gambling, perhaps to individuals and families close to the poverty level, and perhaps to younger people: a myriad of real behaviors and contexts with real welfare consequences. We now have to systematically apply rigorous methods to those settings.

Table 1: Maximum Likelihood Estimates of Discounting Models

Parameter	Point Estimate	Standard Error	*p*-value	95% Confidence Interval	
	A. Exponential Discounting (LL = -18599.6; equation (1))				
δ	0.089	0.008	<0.001	0.074	0.104
r	0.651	0.038	<0.001	0.576	0.726
	B. Quasi-Hyperbolic Discounting (LL = -18596.3; equations (2a) & (2b))				
β	1.003	0.005	<0.001	0.992	1.014
δ	0.073	0.007	<0.001	0.060	0.087
r	0.651	0.038	<0.001	0.577	0.726
	H_0: $\beta = 1$, *p*-value = 0.55				
	C. Fixed Cost Hyperbolic Discounting (LL = -18579.4; equations (3a) & (3b))				
θ	14.353	7.041	0.042	0.552	28.152
β	1.005	0.017	<0.001	0.971	1.038
δ	0.164	0.070	0.018	0.028	0.300
b	-0.014	0.037	0.711	-0.086	0.059
r	0.651	0.038	<0.001	0.577	0.725
	H_0: $\beta = 1$, *p*-value = 0.75; H_0: $\beta = 1$ & b=0, *p*-value = 0.41				
	D. Simple Hyperbolic Discounting (LL = -18598.3; equation (4))				
K	0.089	0.007	<0.001	0.073	0.103
r	0.651	0.038	<0.001	0.576	0.726
	E. General Hyperbolic Discounting (LL = -18596.0; equation (5))				
α	0.497	0.653	0.446	-0.783	1.778
β	0.102	0.024	<0.001	0.055	0.149
r	0.651	0.038	<0.001	0.576	0.726
	F. Weibull Discounting (LL = -18599.0; equation (6))				
ŕ	0.085	0.007	<0.001	0.071	0.097
ś	1.048	0.140	<0.001	0.777	1.318
r	0.651	0.038	<0.001	0.576	0.726
	H_0: ś = 1, *p*-value = 0.73				

Table 2: Estimates of the Effects of Treatments

Parameter	Point Estimate	Standard Error	*p*-value	95% Confidence Interval	
		A. Exponential Discounting (LL = -18459.1; equation (1))			
δ Constant	0.117	0.024	<0.001	0.070	0.164
FED	0.027	0.017	0.126	-0.007	0.061
INFO	-0.036	0.015	0.018	-0.066	-0.006
H_ORDER	-0.037	0.017	0.029	-0.070	-0.004
P_HIGH	-0.001	0.017	0.952	-0.034	0.032
RA_FIRST	0.007	0.020	0.730	-0.032	0.046
FEE_HIGH	-0.014	0.019	0.441	-0.051	0.022
r Constant	0.584	0.062	<0.001	0.461	0.706
RA_FIRST	0.047	0.066	0.476	-0.083	0.177
FEE_HIGH	0.084	0.064	0.190	-0.042	0.210
		B. Quasi-Hyperbolic Discounting (LL = -18389.4; equations (2a) & (2b))			
β Constant	1.003	0.014	<0.001	0.976	1.030
INFO	0.013	0.009	0.170	-0.005	0.025
H_ORDER	0.005	0.010	0.607	-0.015	0.025
P_HIGH	-0.005	0.012	0.968	-0.023	0.022
RA_FIRST	-0.023	0.012	0.051	-0.046	0.001
FEE_HIGH	-0.007	0.010	0.498	-0.027	0.013
δ Constant	0.091	0.020	<0.001	0.051	0.131
FED	0.035	0.017	0.043	0.001	0.069
INFO	-0.021	0.013	0.112	-0.047	0.005
H_ORDER	-0.027	0.016	0.087	-0.058	0.004
P_HIGH	-0.002	0.015	0.873	-0.033	0.028
RA_FIRST	-0.010	0.017	0.578	-0.044	0.024
FEE_HIGH	-0.017	0.016	0.297	-0.049	0.015
r Constant	0.592	0.063	<0.001	0.469	0.716
RA_FIRST	0.038	0.068	0.571	-0.094	0.171
FEE_HIGH	0.084	0.064	0.192	-0.042	0.210

Table 3: Review of Experimental Literature with Real Incentives

Study	Sample (Size)	Elicitation Method	Horizon(s)	Front End Delay(s)	Correct for Non-Linear Utility	Models (estimated rates)	Statistical Method	Hyperbolicky Discounting?
Ainslie & Haendel [1983]	Patients (N=18, 66 choices)	Choice	3 days	0 or 7 days	No	None (n/a)	Counting	Yes
Horowitz [1991]	Students (N=70)	Bidding	64 and 34 days	None	Yes	None (830%, 271%)	Summary statistic	Yes
Kirby & Marakovic [1996]	Students (N=621)	Choice	10 - 70 days	None	No	E, H2 (128% to 1.2E+13)	Non-linear least squares	Yes
Kirby [1997]	Students & Others (N=24, 28, 20)	Bidding	1 - 29 days	None	No	E, H2 (1.6E+04 to 4.3E+10)	Non-linear least squares	Yes
Coller & Williams [1999]	Students (N=199)	Choice	2 months	0 or 1 month	No	E (72.5% with no FED; 30.2% with FED)	Interval regression	Yes (see our Appendix B)
Kirby, Petry & Bickel [1999]		Choice			No			
Anderhub, Güth, Gneezy & Sonsino [2001]	Students	Pricing	4, 8 weeks	None	No	E and H (128% up to 1,084%)	Non-parametric tests	No (see our Appendix B)
Harrison, Lau & Williams [2002]	Danish adults (N=268)	Choice	6, 12, 24 and 36 months	1 month	No	E (28.1%)		No
Kirby & Santiesteban [2003]	Students	Bidding	1-43 days	0 (Experiment 1); 1 day (Experiment 2)	"Not really"	E, H2 (11.3% to 2,877%; 18% to 71,231%)		
Eckel, Johnson & Montmarquette [2005]	Canadian adults	Choice	2 to 28 days	0, 1 day, 2 days, or 2 weeks	No	E		
Harrison, Lau, Rutström & Sullivan [2005]	Danish adults (N=243)	Choice	1, 4, 6, 12, 18, 24 months	1 month	No	E (23.8%)	Interval regression	No
Andersen, Harrison, Lau & Rutström [2008a]	Danish adults (N=243)	Choice	1, 4, 6, 12, 18, 24 months	1 month	Yes	E, H3, W (10.1%)	ML structural estimates	No
Engle-Warnick, Héroux and Montmarquette [2009]	Students (N=151)	Bidding	0, 8, 25 weeks	None	Yes	QH (38% and 33%)	Non-linear least squares	No
Andersen, Harrison, Lau & Rutström [2010]	Students (N=90)	Choice	1, 4 and 6 months	1 month	No	E (27.9%)	Interval regression	No
Dohmen, Falk, Huffman & Sunde [2010]	German adults (N=500)	Choivr	12 months	None	No	None	Interval regression	
Takeuchi [2011]	Students (N=56)	Bidding	Elicited	None	"Not really"	E (726%)	Non-linear least squares	Yes
Benhabib, Bisin & Schotter [2010]	Students (N=27)	Matching	3 days, 1 & 2 weeks, 1,3 & 6 months	None	No	E and FC (≈ 472%)	Non-linear lest squares	Yes
Coller, Harrison & Rutström [2012]	Students (N=87)	Choice	1-60 days	None	Yes	E and QH (mixture) (⋗1000% to 33%)	ML structural estimates	Yes
Andreoni and Sprenger [2012a]	Students (N=97)	Portfolio allocation	35, 70 and 98days	0, 7 and 35 days	Yes	E (30% overall; 28% with no FED)	Non-linear least squares	No
Laury, McInnes and Swarthout [2012]	Students (N=103)	Choice	9 weeks	3 weeks	Yes	E (12.2% and 14.1%)	ML structural estimates	No

Table 4: Estimates of Discounting Models with Student Sample

Parameter	Point Estimate	Standard Error	*p*-value	95% Confidence Interval	
			A. Exponential Discounting (LL = -3441.9; equation (1))		
δ	0.104	0.014	<0.001	0.076	0.131
r	0.534	0.056	<0.001	0.424	0.644
			B. Quasi-Hyperbolic Discounting (LL = -3427.9; equations (2a) & (2b))		
β	0.986	0.010	<0.001	0.966	1.006
δ	0.072	0.012	<0.001	0.048	0.096
r	0.532	0.056	<0.001	0.423	0.642
			H_0: β = 1, *p*-value = 0.18		
			C. Fixed Cost Hyperbolic Discounting (LL = -3418.9; equations (3a) & (3b))		
θ	10.167	26.28	0.70	-41.34	61.67
β	0.966	0.031	<0.001	0.905	1.026
δ	0.123	0.168	0.448	-0.202	0.458
b	-0.060	0.063	0.340	-0.185	0.064
r	0.529	0.056	<0.001	0.421	0.638
			H_0: β = 1, *p*-value = 0.27; H_0: β = 1 & b=0, *p*-value = 0.53		
			D. Simple Hyperbolic Discounting (LL = -3440.6; equation (4))		
K	0.103	0.014	<0.001	0.076	0.131
r	0.534	0.056	<0.001	0.424	0.643
			E. Generalized Hyperbolic (LL = -3427.9; equation (5))		
α	5.540	7.53	0.462	-9.22	20.30
β	0.261	0.191	0.173	-0.114	0.636
r	0.532	0.056	<0.001	0.423	0.642
			F. Weibull Discounting (LL = -3425.9; equation (6))		
ŕ	0.094	0.012	<0.001	0.070	0.118
ś	1.59	0.428	<0.001	0.758	2.430
r	0.532	0.056	<0.001	0.423	0.642
			H_0: ś = 1, *p*-value = 0.17		

Figure 1: Estimated Distribution of δ Assuming the Exponential Discounting Model

Estimates generated with different econometric specifications for heterogeneity

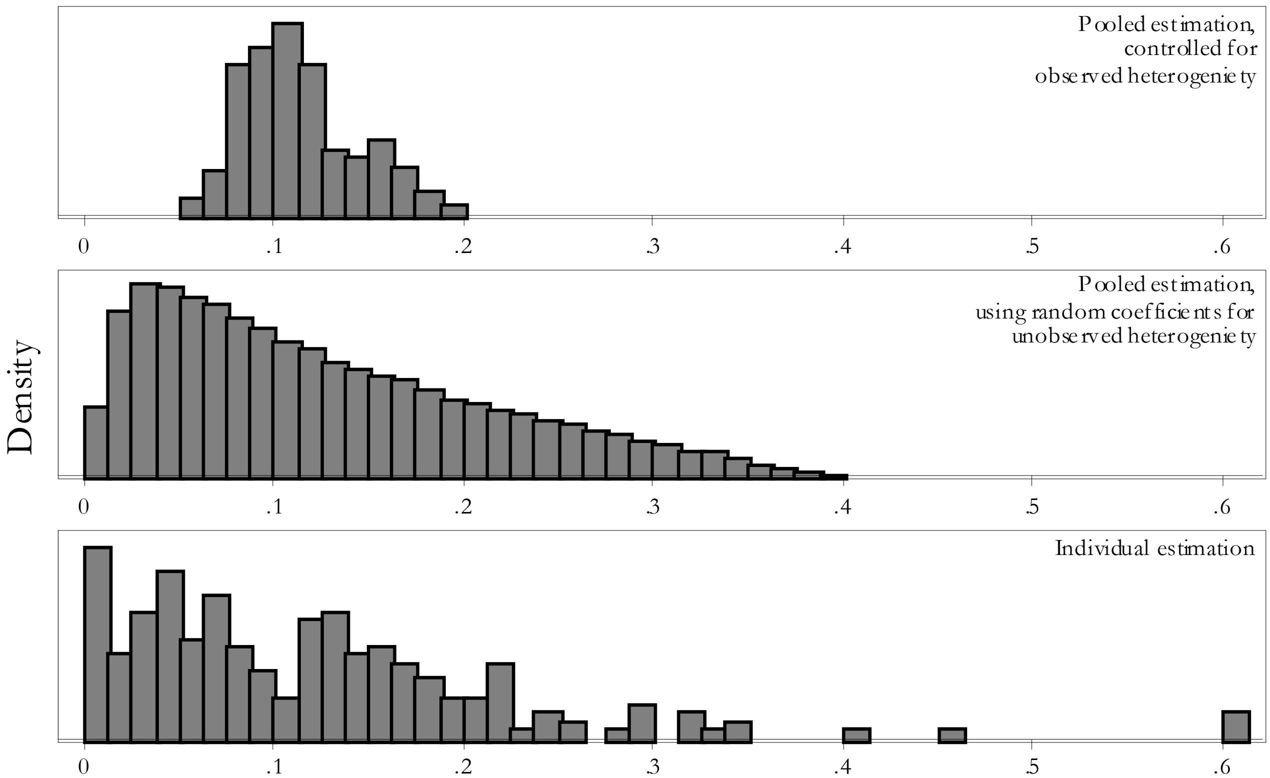

Figure 2: Estimated Distributions of β Assuming the Quasi-Hyberbolic Discounting Model

Estimates generated with different econometric specifications for heterogeneity

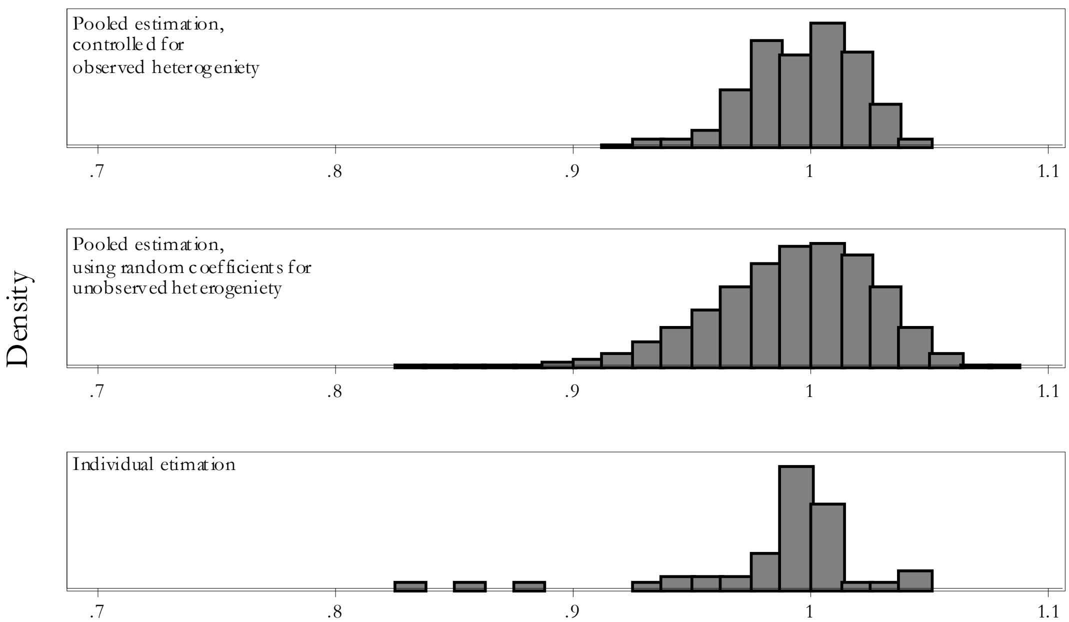

Figure 3: Estimated Distributions of δ Assuming the Quasi-Hyberbolic Discounting Model

Estimates generated with different econometric specifications for heterogeneity

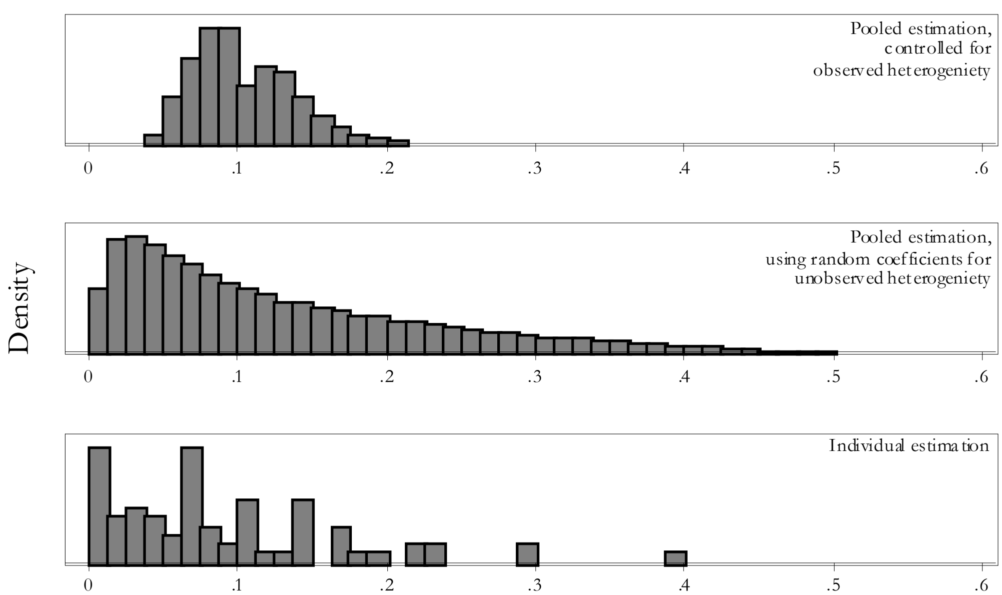

References

Ainslie, George, "Specious Reward: A Behavioral Theory of Impulsiveness and Impulse Control," *Psychological Bulletin*, 92(4), 1975, 463-496.

Ainslie, George, *Picoeconomics* (New York: Cambridge University Press, 1992).

Ainslie, George, and Haendel, V., "The Motives of the Will," in E. Gottheil, K. Druley, T. Skodola, and H. Waxman (eds.), *Etiology Aspects of Alcohol and Drug Abuse* (Springfield, IL: Charles C. Thomas, 1983).

Ahlbrecht, Martin, and Weber, Martin, "An Empirical Study on Intertemporal Decision Making under Risk," *Management Science*, 43(6), June 1997, 813-826.

Anderhub, Vital; Güth, Werner; Gneezy, Uri, and Sonsino, Doron, "On the Interaction of Risk and Time Preference: An Experimental Study," *German Economic Review*, 2(3), 2001, 239-253.

Andersen, Steffen; Cox, James C.; Harrison, Glenn W.; Lau, Morten; Rutström, E. Elisabet, and Sadiraj, Vjollca, "Asset Integration and Attitudes to Risk: Theory and Evidence," *Working Paper 2011-17*, Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University, 2011.

Andersen, Steffen; Harrison, Glenn W.; Lau, Morten I., and Rutström, E. Elisabet, "Eliciting Risk and Time Preferences," *Econometrica*, 76(3), 2008a, 583-619.

Andersen, Steffen, Harrison, Glenn W.; Lau, Morten I., and Rutström, E. Elisabet, "Lost in State Space: Are Preferences Stable?" *International Economic Review*, 49(3), 2008b, 1091-1112.

Andersen, Steffen, Harrison, Glenn W.; Lau, Morten I., and Rutström, E. Elisabet, "Preference Heterogeneity in Experiments: Comparing the Field and Laboratory," *Journal of Economic Behavior and Organization*, 73, 2010, 209-224.

Andersen, Steffen; Harrison, Glenn W.; Lau, Morten, and Rutström, E. Elisabet, "Multiattribute Utility Theory, Intertemporal Utility and Correlation Aversion," *Working Paper 2011-04*, Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University, 2011.

Andersen, Steffen; Harrison, Glenn W., Lau, Morten I., and Rutström, E. Elisabet, "Discounting Behavior and the Magnitude Effect," *Economica*, 80, 2013, 670-697.

Andreoni, James, and Sprenger, Charles, "Estimating Time Preferences from Convex Budgets," *American Economic Review*, 102(7), December 2012a, 3333-3356.

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

Benhabib, Jess; Bisin, Alberto, and Schotter, Andrew, "Present-Bias, Quasi-Hyperbolic Discounting, and Fixed Costs," *Games & Economic Behavior*, 69(2), July 2010, 205-223.

Camerer, Colin, and Ho, Teck-Hua, "Violations of the Betweenness Axiom and Nonlinearity in Probability," *Journal of Risk & Uncertainty*, 8, 1994, 167-196.

Coller, Maribeth; Harrison, Glenn W., and Rutström, E. Elisabet, "Latent Process Heterogeneity in Discounting Behavior," *Oxford Economic Papers*, 64, 2012, 375-391.

Coller, Maribeth, and Williams, Melonie B, "Eliciting Individual Discount Rates," *Experimental Economics*, 2, 1999, 107-127.

Cubitt, Robin P., and Read, Daniel, "Can Intertemporal Choice Experiments Elicit Time Preferences for Consumption?" *Experimental Economics*, 10, 2007, 369-389.

Dohmen, Thomas; Falk, Armin; Huffman, David, and Sunde, Uwe, "Are Risk Aversion and Impatience Related to Cognitive Ability?," *American Economic Review*, 100(3), 2010, 1238-1260.

Ebert, Jane E.J., and Prelec, Drazen, "The Fragility of Time: Time-Insensitivity and Valuation of the Near and Far Future," *Management Science*, 53(9), 2007, 1423-1438.

Eckel, Catherine; Johnson, Cathleen, and Montmarquette, Claude, "Savings Decisions of the Working Poor: Short- and Long-term Horizons," in J. Carpenter, G. W. Harrison and J. A. List (eds.), Field Experiments in Economics (Greenwich, CT: JAI Press, *Research in Experimental Economics*, Volume 10, 2005).

Elster, Jon, *Ulysses and the Sirens: Studies in Rationality and Irrationality* (Cambridge University Press, Cambridge, 1979).

Engle-Warnick, Jim; Héroux, Julie, and Montmarquette, Claude, "Willingness to Pay to Reduce Future Risk," *CIRANO Scientific Series 2009s-27*, Montreal, Canada, 2009.

Epper, Thomas; Fehr-Duda, Helga, and Bruhin, Adrian, "Viewing the Future through a Warped Lens: Why Uncertainty Generates Hyperbolic Discounting," *Working Paper*, ETH Zurich, September 2010.

Frederick, Shane; Loewenstein, George; and O'Donoghue, Ted, "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature*, 40, June 2002, 351-401.

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

Halevy, Yoram, "Strotz Meets Allais: Diminishing Impatience and the Certainty Effect," *American Economic Review*, 98(3), 2008, 1145-1162.

Harrison, Glenn W., "Theory and Misbehavior of First-Price Auctions: Reply," *American Economic Review*, 82, December 1992, 1426-1443.

Harrison, Glenn W., "Field Experiments and Control," in J. Carpenter, G.W. Harrison and J.A. List (eds.), *Field Experiments in Economics* (Greenwich, CT: JAI Press, Research in Experimental

Economics, Volume 10, 2005).

Harrison, Glenn W., "Hypothetical Bias Over Uncertain Outcomes," in J.A. List (ed)., *Using Experimental Methods in Environmental and Resource Economics* (Northampton, MA: Elgar, 2006).

Harrison, Glenn W.; Johnson, Eric; McInnes, Melayne M., and Rutström, E. Elisabet, "Risk Aversion and Incentive Effects: Comment," *American Economic Review*, 95(3), June 2005, 897-901.

Harrison, Glenn W.; Lau, Morten, and Rutström, Elisabet, "Estimating Risk Attitudes in Denmark: A Field Experiment," *Scandinavian Journal of Economics*, 109(2), June 2007, 341-368.

Harrison, Glenn W.; Lau, Morten I.; Rutström, E. Elisabet, and Sullivan, Melonie B., "Eliciting Risk and Time Preferences Using Field Experiments: Some Methodological Issues," in J. Carpenter, G.W. Harrison and J.A. List (eds.), *Field Experiments in Economics* (Greenwich, CT: JAI Press, Research in Experimental Economics, Volume 10, 2005).

Harrison, Glenn W.; Lau, Morten I., and Williams, Melonie B., "Estimating Individual Discount Rates for Denmark: A Field Experiment," *American Economic Review*, 92(5), December 2002, 1606-1617.

Harrison, Glenn W., and List, John A., "Field Experiments," *Journal of Economic Literature*, 42(4), December 2004, 1013-1059.

Harrison, Glenn W., and Rutström, E. Elisabet, "Risk Aversion in the Laboratory," in J.C. Cox and G.W. Harrison (eds.), *Risk Aversion in Experiments* (Bingley, UK: Emerald, Research in Experimental Economics, Volume 12, 2008a).

Harrison, Glenn W., and Rutström, E. Elisabet,"Experimental Evidence on the Existence of Hypothetical Bias in Value Elicitation Experiments," in C.R. Plott and V.L. Smith (eds.), *Handbook of Experimental Economics Results* (New York: Elsevier Press, 2008b).

Harstad, Ronald M., "Dominant Strategy Adoption and Bidders' Experience with Pricing Rules," *Experimental Economics*, 3(3), December 2000, 261-280.

Hey, John D., and Orme, Chris, "Investigating Generalizations of Expected Utility Theory Using Experimental Data," *Econometrica*, 62(6), November 1994, 1291-1326.

Holden, Stein T.; Shiferaw, Bekele; Wik, Mette, "Poverty, Market Imperfections and Time Preferences of Relevance for Environmental Policy?" *Environment and Development Economics*, 3, 1998, 105-130.

Holt, Charles A., and Laury, Susan K., "Risk Aversion and Incentive Effects," *American Economic Review*, 92(5), December 2002, 1644-1655.

Holt, Charles A., and Laury, Susan K., "Risk Aversion and Incentive Effects: New Data Without Order Effects," *American Economic Review*, 95(3), June 2005, 902-912.

Horowitz, John K., "Discounting Money Payoffs: An Experimental Analysis," in S. Kaish and B. Gilad (eds.), *Handbook of Behavioral Economics* (Greenwich: JAI Press, Volume 2B, 1991).

Keller, L. Robin, and Strazzera, Elisabetta, "Examining Predictive Accuracy Among Discounting Models," *Journal of Risk and Uncertainty*, 24(2), March 2002, 143-160.

Keren, Gideon, and Roelofsma, Peter, "Immediacy and Certainty in Intertemporal Choice," *Organizational Behavior and Human Decision Processes*, 63(3), September 1995, 287-297.

Kirby, Kris N., "Bidding on the Future: Evidence Against Normative Discounting of Delayed Rewards," *Journal of Experimental Psychology: General*, 126, 1997, 54-70.

Kirby, Kris N., and Guastello, Barbarose, "Making Choices in Anticipation of Similar Future Choices Can Increase Self-Control," *Journal of Experimental Psychology: Applied*, 7(2), 2001, 154-164.

Kirby, Kris N., and Maraković, Nino N., "Modeling myopic decisions: Evidence of hyperbolic delay-discounting within subjects and amounts," *Organizational Behavior & Human Decision Processes*, 64, 1995, 22-30.

Kirby, Kris N., and Maraković, Nino N., "Delay-discounting probabilistic rewards: Rates decrease as amounts increase," *Psychonomic Bulletin & Review*, 1996, 3:1, 100-104.

Kirby, Kris N.; Petry, Nancy M., and Bickel, Warren, "Heroin Addicts Have Higher Discount Rates for Delayed Rewards than Non-Drug-Using Controls," *Journal of Experimental Psychology: General*, 1999, 128(1), 78-87.

Kirby, Kris N., and Santiesteban, Mariana, "Concave Utility, Transactions Costs, and Risk in Measuring Discounting of Delayed Rewards," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(1), 2003, 66-79.

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

Laury, Susan K.; McInnes, Melayne Morgan, and Swarthout, J. Todd, "Avoiding the Curves: Direct Elicitation of Time Preferences," *Journal of Risk and Uncertainty*, 44, 2012, 181-217.

Loewenstein, George; O'Donoghue, Ted, and Rabin, Matthew, "Projection Bias in Predicting Future Utility," *Quarterly Journal of Economics*, 118(4), November 2003, 1209-1248.

Loewenstein, George, and Prelec, Drazen, "Anomalies in Intertemporal Choice: Evidence and Interpretation," *Quarterly Journal of Economics*, 107, 1992, 573-592.

Mazur, James E., "Tests of an Equivalence Rule for Fixed and Variable Reinforcer Delays," *Journal of Experimental Psychology: Animal Behavior Processes*, 10(4), 1984, 426-437.

Phelps, Edmund S., and Pollak, Robert A., "On Second-Best National Saving and Game-Equilibrium Growth," *Review of Economic Studies*, 35, 1968, 185-199.

Prelec, Drazen, "Decreasing Impatience: A Criterion for Non-Stationary Time Preference and 'Hyperbolic' Discounting," *Scandinavian Journal of Economics*, 106(3), 2004, 511-532.

Read, Daniel, "Is Time-Discounting Hyperbolic or Subadditive?" *Journal of Risk and Uncertainty*, 23(1), 2001, 5-32.

Read, Daniel; Frederick, Shane; Orsel, Burcu, and Rahman, Juwaria, "Four Score and Seven Years from Now: The Date/Delay Effect in Temporal Discounting," *Management Science*, 51(9), September 2005, 1326-1335.

Read, Daniel, and Roelofsma, Peter H.M.P., "Subadditive versus Hyperbolic Discounting: A Comparison of Choice and Matching," *Organizational Behavior and Human Decision Processes*, 91, 2003, 140-153.

Rutström, E. Elisabet, "Home-Grown Values and the Design of Incentive Compatible Auctions," *International Journal of Game Theory*, 27(3), 1998, 427-441.

Takeuchi, Kan, "Non-Parametric Test of Time Consistency: Present Bias and Future Bias," *Games & Economic Behavior*, 71(2), 2011, 456-478.

Warner, John T., and Pleeter, Saul, "The Personal Discount Rate: Evidence from Military Downsizing Programs," *American Economic Review*, 91(1), March 2001, 33-53.

Weber, Bethany, and Chapman, Gretchen 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, 2005, 104-118.

Appendix A: Instructions (WORKING PAPER)

We document the instructions by first listing the "manuscript" that shows what was given to subjects and read to them, and then we document some of the screen displays. The original Danish manuscript is available on request. The originals were in 14-point font, printed on A4 paper for nice page breaks (a horizontal line below indicates a page break), and given to subjects in laminated form. The manuscript below was for the sessions in which the discount rate task was presented first. After these experimental tasks were completed there were additional tasks in the session that are not relevant here.

A. Experimental Manuscript

Welcome announcement

[Give informed consent form to subjects.]

Thank you for agreeing to participate in this survey. The survey is financed by the Social Science Research Council and the Carlsberg Foundation and concerns the economics of decision making.

Before we begin the survey, let me read out the informed consent form that is handed out to you. This form explains your rights as a participant in the survey, what the survey is about and how we make payments to you.

[Read the informed consent form.]

Is everyone able to stay for the full two hours of the meeting? Before we begin, I will ask each of you to pick an envelope from me. The envelope contains a card with an ID number that we will use to keep track of who answered which questions. All records and published results will be linked to anonymous ID numbers only, and not to your name. Please keep your ID numbers private and do not share the information with anyone else.

[Each subject picks an envelope.]

You will be given written instructions during the survey, but make all decisions on the computer in front of you. Please enter your ID number on the computer in front of you, but keep the card for later use.

You will now continue with the first task. The problem is not designed to test you. The only right answer is what you really would choose. That is why the task gives you the chance of winning money. I will now distribute the instructions and then read it out loud.

[Give IDR instructions to subjects.]
[Read the IDR instructions.]

Task D

In this task you will make a number of choices between two options labeled "A" and "B". An example of your task is shown on the right. You will make all decisions on a computer.

All decisions have the same format. In the example on the right Option A pays 100 kroner today and Option B pays 105 kroner twelve months from now. By choosing option B you would get an annual return of 5% on the 100 kroner.

We will present you with 40 of these decisions. The only difference between them is that the amounts and payment dates in Option A and B will differ.

You will have a 1-in-10 chance of being paid for one of these decisions. The selection is made with a 10-sided die. If the roll of the die gives the number 1 you will be paid for one of the 40 decisions, but if the roll gives any other number you will not be paid. If you are paid for one of these 40 decisions, then we will further select one of these decisions by rolling a 4-sided and a 10-sided die. When you make your choices you will not know which decision is selected for payment. You should therefore treat each decision as if it might actually count for payment.

You will receive the money on the date stated in your preferred option. If you receive some money today, then it is paid out at the end of the experiment. If you receive some money to be paid in the future, then it is transferred to your personal bank account on the specified date. In that case you will receive a written confirmation from Copenhagen Business School which guarantees that the money is reserved on an account at Danske Bank. You can send this document to Danske Bank in a prepaid envelope, and the bank will transfer the money to your account on the specified date.

Before making your choices you will have a chance to practice so that you better understand the consequences of your choices. Please proceed on the computer to the practice task. You will be paid in caramels for this practice task, and they are being paid on the time stated in your preferred option.

[Subjects make decisions in the practice IDR task.]

I will now come around and pay you in caramels for your choice of A or B. Please proceed to the actual task after your earnings are recorded. You will have a 1-in-10 chance of being paid for one of the 40 decisions in the actual task.

Password 1:____

[Subjects make decisions in the actual IDR task.]

I will now come around and ask you to roll a 10-sided die to determine if you are being paid for one of the decisions. If the roll of the die gives the number 1 you will be paid for one of the 40 decisions, but if the roll gives any other number you will not be paid. If you are paid for one of the 40 decisions, then I will ask you to roll a 4-sided and a 10-sided die to select one of the decisions for payment.

Password 2:____

[Roll 10-sided die to determine if they are being paid.]
[Roll 4-sided and 10-sided dice to determine the decision for payment.]

You will now continue with the second task. I will distribute the instructions and then read it out loud.

[Give RA instructions to subjects.]
[Read the RA instructions.]

Task L

In this task you will make a number of choices between two options labeled "A" and "B". An example of your task is shown on the right. You will make all decisions on a computer.

All decisions have the same format. In the example on the right Option A pays 60 kroner if the outcome of a roll of a ten-sided die is 1, and it pays 40 kroner if the outcome is 2-10. Option B pays 90 kroner if the outcome of the roll of the die is 1 and 10 kroner if the outcome is 2-10. All payments in this task are made today at the end of the experiment.

We will present you with 40 such decisions. The only difference between them is that the probabilities and amounts in Option A and B will differ.

You have a 1-in-10 chance of being paid for one of these decisions. The selection is made with a 10-sided die. If the roll of the die gives the number 1 you will be paid for one of the 40 decisions, but if the roll gives any other number you will not be paid. If you are paid for one of these 40 decisions, then we will further select one of these decisions by rolling a 4-sided and a 10-sided die. A third die roll with a 10-sided die determines the payment for your choice of Option A or B. When you make your choices you will not know which decision is selected for payment. You should therefore treat each decision as if it might actually count for payment.

If you are being paid for one of the decisions, we will pay you according to your choice in the selected decision. You will then receive the money at the end of the experiment.

Before making your choices you will have a chance to practice so that you better understand the consequences of your choices. Please proceed on the computer to the practice task. You will be paid in caramels for this practice task.

[Subjects make decisions in the practice RA task.]

I will now come around and pay you in caramels for your choice of A or B. I will ask you to roll a 10-sided die to determine the payment for your choice of A or B. Please proceed to the actual task after your earnings are recorded. You will have a 1-in-10 chance of being paid for one of the 40 decisions in the actual task.

Password 3:____

[Subjects make decisions in the actual RA task.]

I will now come around and ask you to roll a 10-sided die to determine if you are being paid for one of the decisions. If the roll of the die gives the number 1 you will be paid for one of the 40 decisions, but if the roll gives any other number you will not be paid. If you are paid for one of the 40 decisions, then I will ask you to roll a 4-sided and a 10-sided die to select one of the decisions for payment. A third die roll with a 10-sided die determines the payment for your choice of Option A or B.

Password 4:____

[Roll 10-sided die to determine if they are being paid.]
[Roll 4-sided and 10-sided dice to determine the decision for payment.]
[Roll 10-sided die to determine payment in Option A and B.]

You will now continue with the third task. I will distribute the instructions and then read it out loud.

[ADDITIONAL INSTRUCTIONS WERE PROVIDED HERE]

B. Typical Screen Shots for Lottery Choices

The first screen shot on the next page shows the full screen within which the text box is contained, so that one gets an impression of what the subject encountered in all screen shots. Then we display more detailed screen shots of the practice example and the first few lottery choices. Prior to each block of 10 lottery choices the subject was told that the lottery prizes for the next 10 choices would stay the same and the only thing that would vary would be the probabilities. We then show the sequence of the first two lotteries, and then lottery 11 which uses new prizes.

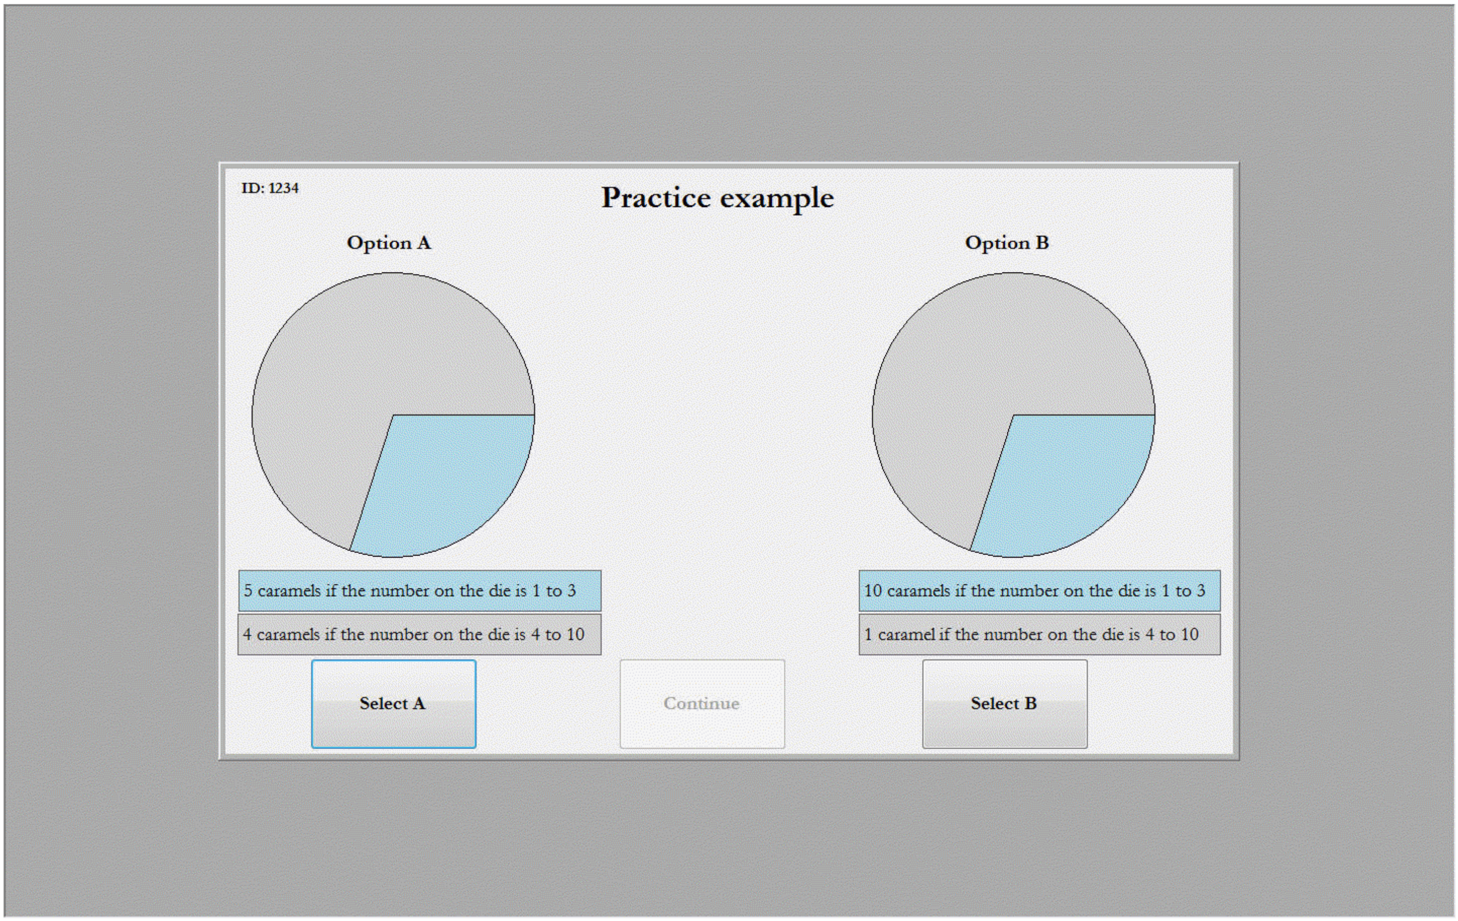
ID: 1234
Practice example
Option A
Option B
5 caramels if the number on the die is 1 to 3
4 caramels if the number on the die is 4 to 10
10 caramels if the number on the die is 1 to 3
1 caramel if the number on the die is 4 to 10
Select A
Continue
Select B

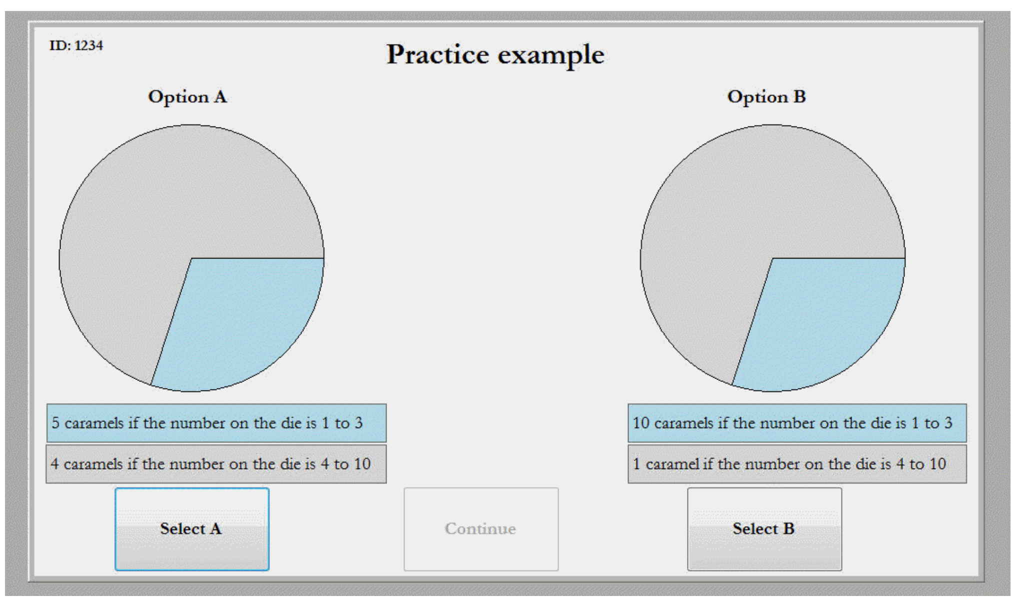
ID: 1234
Practice example
Option A
Option B
5 caramels if the number on the die is 1 to 3
4 caramels if the number on the die is 4 to 10
10 caramels if the number on the die is 1 to 3
1 caramel if the number on the die is 4 to 10
Select A
Continue
Select B

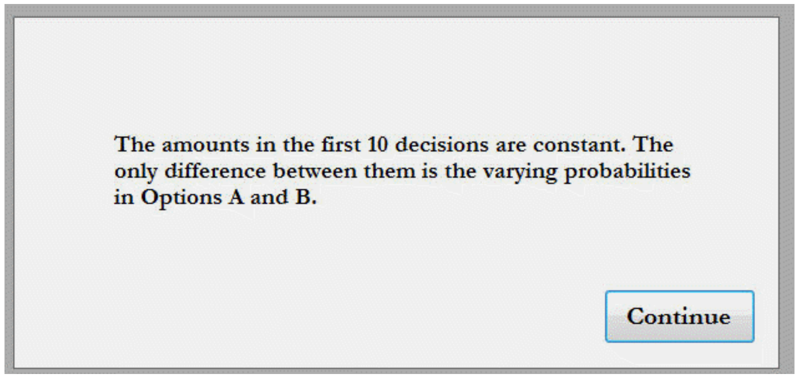
The amounts in the first 10 decisions are constant. The only difference between them is the varying probabilities in Options A and B.
Continue

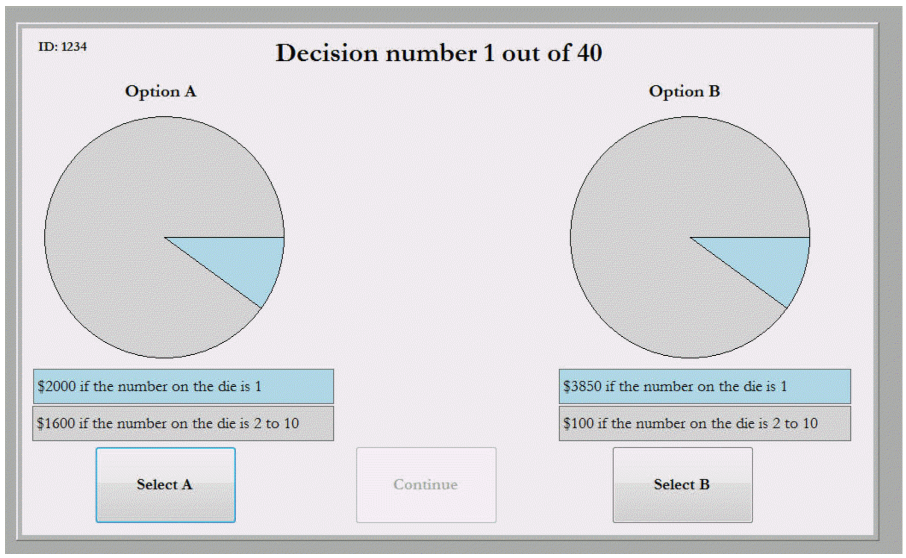
ID: 1234
Decision number 1 out of 40
Option A
Option B
$2000 if the number on the die is 1
$1600 if the number on the die is 2 to 10
$3850 if the number on the die is 1
$100 if the number on the die is 2 to 10
Select A
Continue
Select B

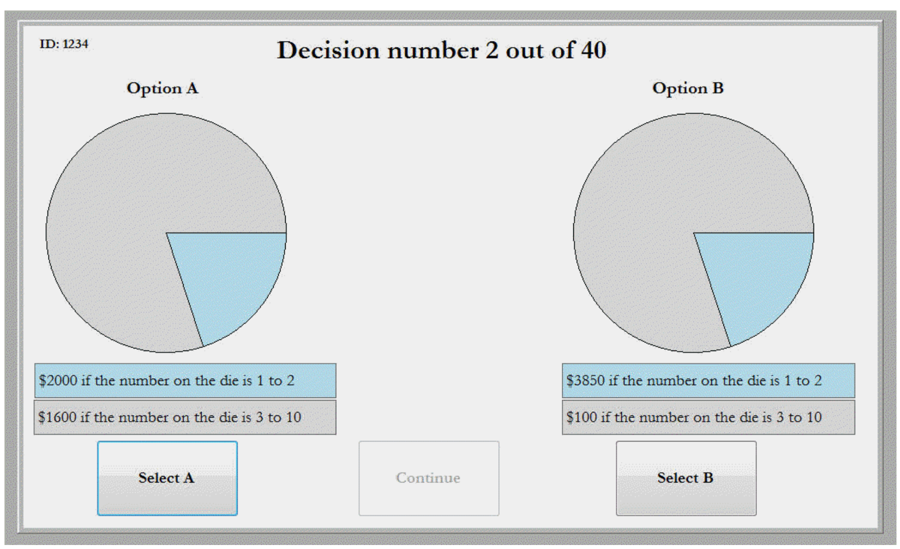
ID: 1234
Decision number 2 out of 40
Option A
Option B
$2000 if the number on the die is 1 to 2
$1600 if the number on the die is 3 to 10
$3850 if the number on the die is 1 to 2
$100 if the number on the die is 3 to 10
Select A
Continue
Select B

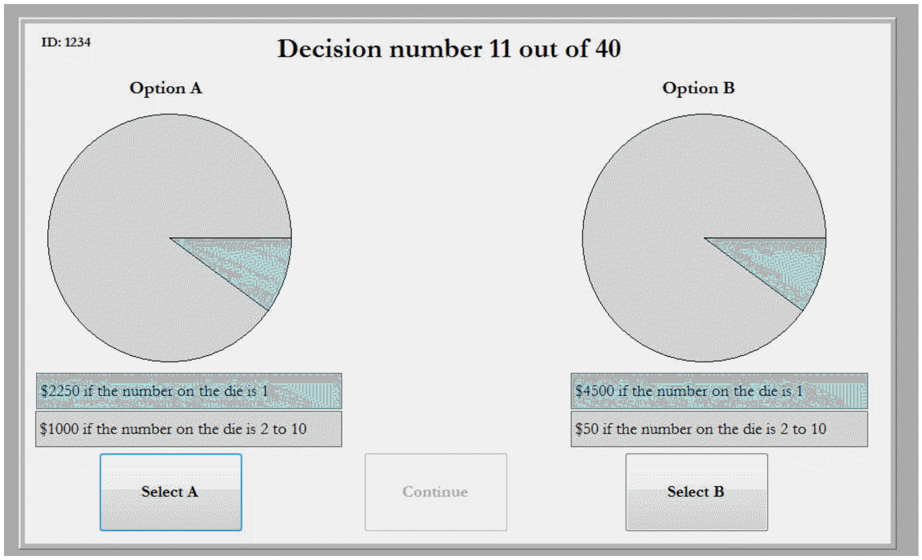

C. Typical Screen Shots for Discounting Choices

The next page shows the practice example provided at the beginning of these tasks. The top panel shows the initial screen shot, and then the next two panels show how the selected option is highlighted to make it clear to the subject which option is being selected.

The following page shows the information that was given to each subject prior to each block of 10 choices. This information was that the principal and horizon would remain constant for the next 10 choices, but that the only thing that would change would be the amount in the "later" option. In these displays the implied interest rate is displayed.

Finally, after the first 10 choices a new horizon was selected for the next 10 choices.

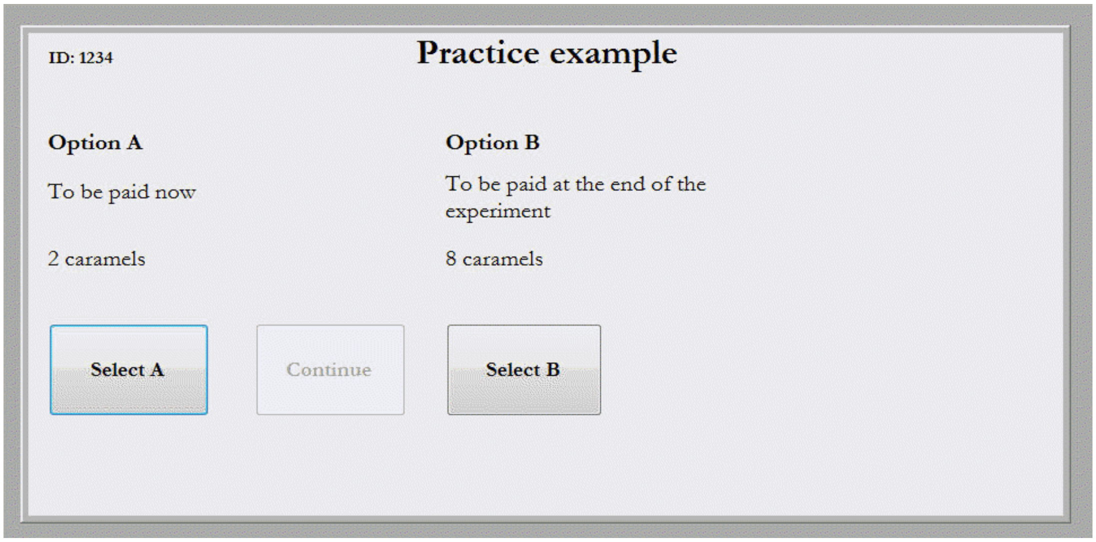
ID: 1234
Practice example
Option A
To be paid now
2 caramels
Option B
To be paid at the end of the experiment
8 caramels
Select A
Continue
Select B

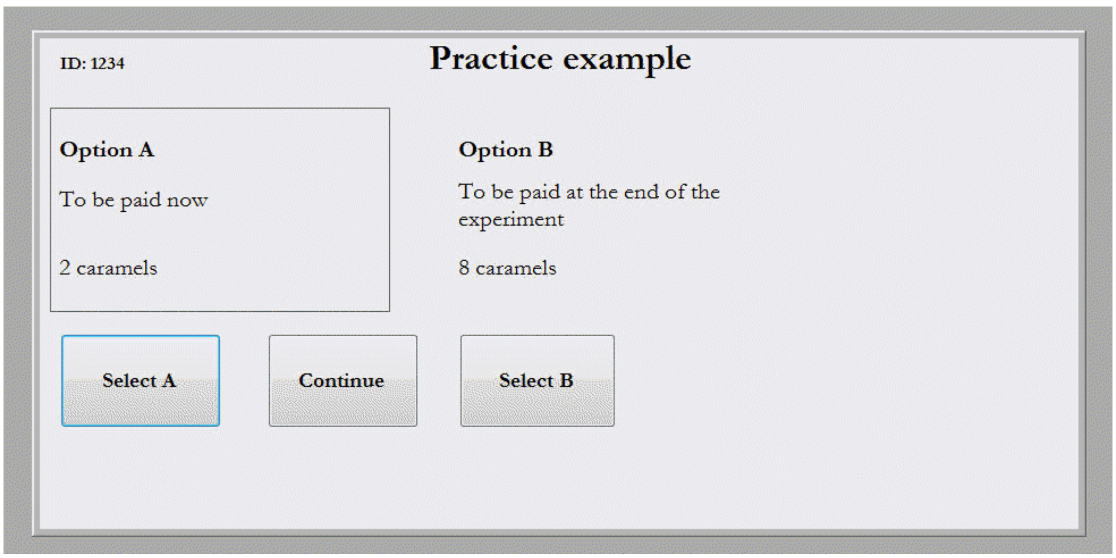
ID: 1234
Practice example
Option A
To be paid now
2 caramels
Option B
To be paid at the end of the experiment
8 caramels
Select A
Continue
Select B

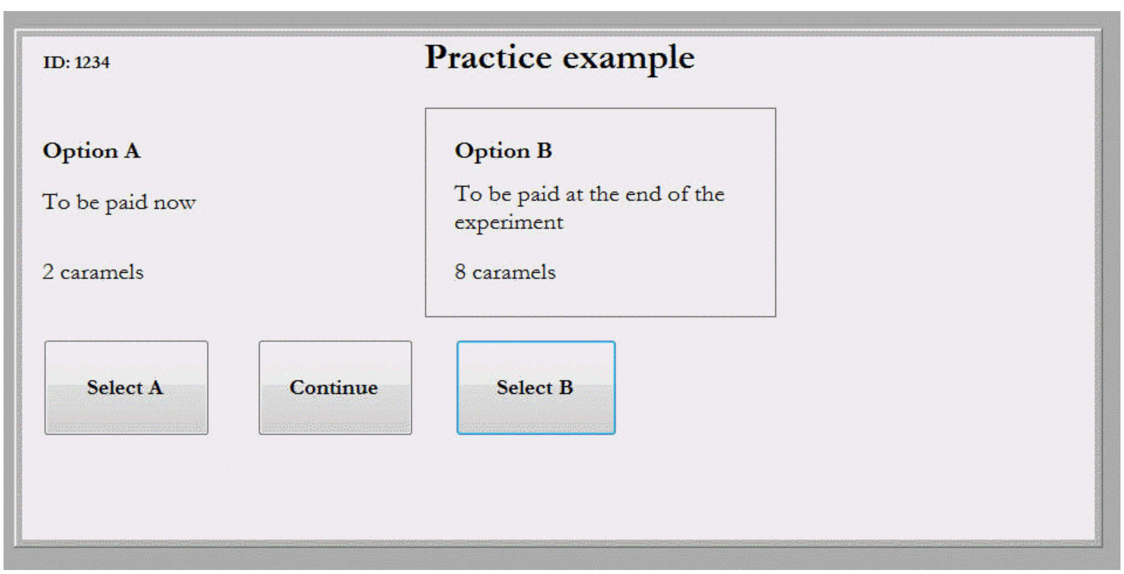
ID: 1234
Practice example
Option A
To be paid now
2 caramels
Option B
To be paid at the end of the experiment
8 caramels
Select A
Continue
Select B

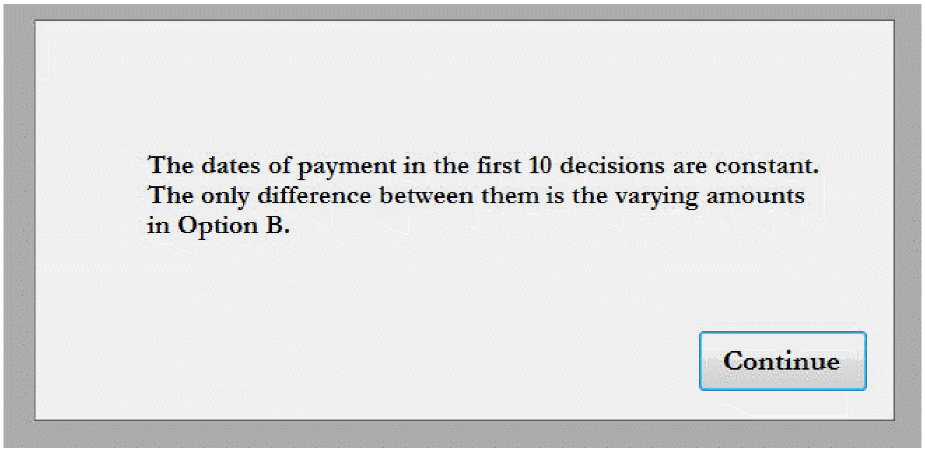
The dates of payment in the first 10 decisions are constant.
The only difference between them is the varying amounts
in Option B.
Continue

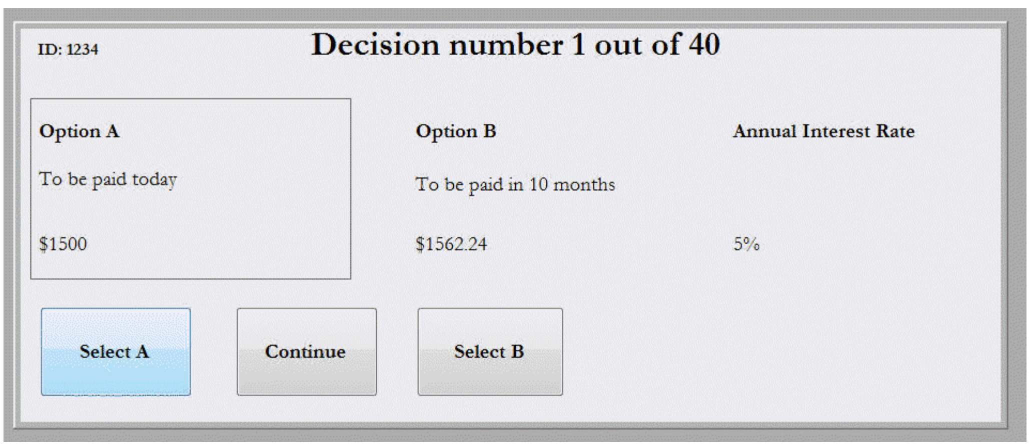
ID: 1234
Decision number 1 out of 40
Option A
To be paid today
$1500
Option B
To be paid in 10 months
$1562.24
Annual Interest Rate
5%
Select A
Continue
Select B

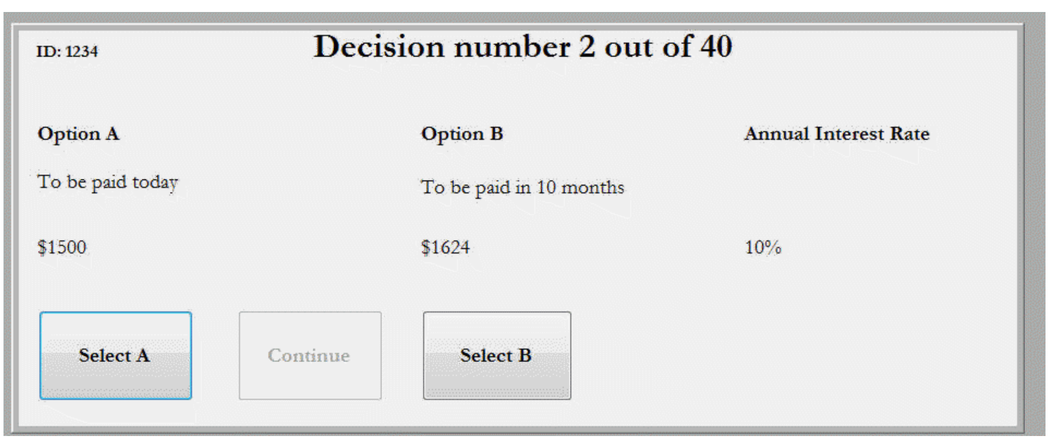
ID: 1234
Decision number 2 out of 40
Option A
To be paid today
$1500
Option B
To be paid in 10 months
$1624
Annual Interest Rate
10%
Select A
Continue
Select B

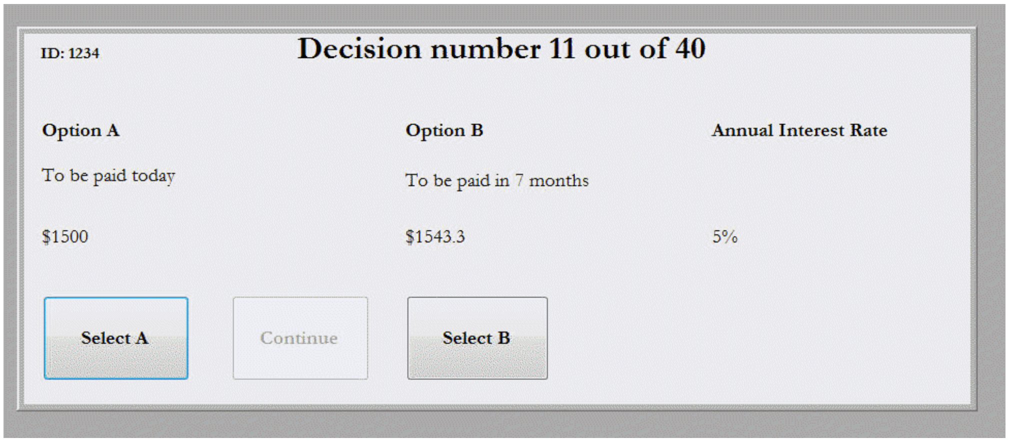

D. Parameter Values

Table A1 shows the parameters of the lottery choice tasks, and Table A2 shows the parameters of the discounting choice tasks.

In Table A1 the parameters are (1) the decision number, (2) the probability of the high prize in each lottery, (3) the high prize of lottery A, in kroner, (4) the low prize of lottery A, in kroner, (5) the high prize of lottery B, in kroner, (6) the low prize of lottery B, in kroner, (7) the expected value of lottery A, and (8) the expected value of lottery B. The information in columns (7) and (8) was not presented to subjects.

Table A1: Parameters for Lottery Choices

Decision	Probability of High Prize	Lottery A High Prize	Lottery A Low Prize	Lottery B High Prize	Lottery B Low Prize	EV of Lottery A	EV of Lottery B
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1	0.1	1125	750	2000	250	787.5	425
2	0.2	1125	750	2000	250	825	600
3	0.3	1125	750	2000	250	862.5	775
4	0.4	1125	750	2000	250	900	950
5	0.5	1125	750	2000	250	937.5	1125
6	0.6	1125	750	2000	250	975	1300
7	0.7	1125	750	2000	250	1012.5	1475
8	0.8	1125	750	2000	250	1050	1650
9	0.9	1125	750	2000	250	1087.5	1825
10	1	1125	750	2000	250	1125	2000
11	0.1	1000	875	2000	75	887.5	267.5
12	0.2	1000	875	2000	75	900	460
13	0.3	1000	875	2000	75	912.5	652.5

14	0.4	1000	875	2000	75	925	845
15	0.5	1000	875	2000	75	937.5	1037.5
16	0.6	1000	875	2000	75	950	1230
17	0.7	1000	875	2000	75	962.5	1422.5
18	0.8	1000	875	2000	75	975	1615
19	0.9	1000	875	2000	75	987.5	1807.5
20	1	1000	875	2000	75	1000	2000
21	0.1	2000	1600	3850	100	1640	475
22	0.2	2000	1600	3850	100	1680	850
23	0.3	2000	1600	3850	100	1720	1225
24	0.4	2000	1600	3850	100	1760	1600
25	0.5	2000	1600	3850	100	1800	1975
26	0.6	2000	1600	3850	100	1840	2350
27	0.7	2000	1600	3850	100	1880	2725
28	0.8	2000	1600	3850	100	1920	3100
29	0.9	2000	1600	3850	100	1960	3475
30	1	2000	1600	3850	100	2000	3850
31	0.1	2250	1000	4500	50	1125	495
32	0.2	2250	1000	4500	50	1250	940
33	0.3	2250	1000	4500	50	1375	1385
34	0.4	2250	1000	4500	50	1500	1830
35	0.5	2250	1000	4500	50	1625	2275
36	0.6	2250	1000	4500	50	1750	2720
37	0.7	2250	1000	4500	50	1875	3165
38	0.8	2250	1000	4500	50	2000	3610
39	0.9	2250	1000	4500	50	2125	4055
40	1	2250	1000	4500	50	2250	4500

In Table A2 the parameters are (1) the horizon in months, (2) the task number in sequence if this horizon was selected for the subject to make choices over, (3) the principal of 3000 kroner if the subject had the "higher stakes" condition, (4) the principal of 1500 kroner if the subject had the "lower stakes" condition, (5) the annual interest rate presented to the subject if that treatment was applied (this is also the annual effective rate with annual compounding), (6) the delayed payment if the subject had the "higher stakes" condition, (7) the delayed payment if the subject had the "lower stakes" condition, (8) the implied annual effective rate with quarterly compounding, and (9) the implied annual effective rate with daily compounding. The values in columns (8) and (9) were not presented to subjects.

Table A2: Parameters for Discounting Choices

Horizon in months	Task	Principal in high stakes	Principal if low stakes	Annual Interest Rate	Delayed Payment if high stakes	Delayed Payment if low stakes	AER Quarterly	AER Daily
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
0.5	1	3000	1500	5%	3006.10	1503.05	5.1%	5.1%
0.5	2	3000	1500	10%	3011.94	1505.97	10.4%	10.5%
0.5	3	3000	1500	15%	3017.52	1508.76	15.9%	16.2%
0.5	4	3000	1500	20%	3022.88	1511.44	21.6%	22.1%
0.5	5	3000	1500	25%	3028.02	1514.01	27.4%	28.4%
0.5	6	3000	1500	30%	3032.98	1516.49	33.5%	35.0%
0.5	7	3000	1500	35%	3037.75	1518.87	39.9%	41.9%
0.5	8	3000	1500	40%	3042.36	1521.18	46.4%	49.1%
0.5	9	3000	1500	45%	3046.81	1523.40	53.2%	56.8%
0.5	10	3000	1500	50%	3051.11	1525.56	60.2%	64.8%
1	1	3000	1500	5%	3012.22	1506.11	5.1%	5.1%
1	2	3000	1500	10%	3023.92	1511.96	10.4%	10.5%
1	3	3000	1500	15%	3035.14	1517.57	15.9%	16.2%
1	4	3000	1500	20%	3045.93	1522.96	21.6%	22.1%
1	5	3000	1500	25%	3056.31	1528.15	27.4%	28.4%
1	6	3000	1500	30%	3066.31	1533.16	33.5%	35.0%
1	7	3000	1500	35%	3075.97	1537.99	39.9%	41.9%
1	8	3000	1500	40%	3085.31	1542.65	46.4%	49.1%
1	9	3000	1500	45%	3094.34	1547.17	53.2%	56.8%
1	10	3000	1500	50%	3103.10	1551.55	60.2%	64.8%
2	1	3000	1500	5%	3024.49	1512.25	5.1%	5.1%
2	2	3000	1500	10%	3048.04	1524.02	10.4%	10.5%
2	3	3000	1500	15%	3070.70	1535.35	15.9%	16.2%
2	4	3000	1500	20%	3092.56	1546.28	21.6%	22.1%
2	5	3000	1500	25%	3113.67	1556.84	27.4%	28.4%
2	6	3000	1500	30%	3134.09	1567.05	33.5%	35.0%
2	7	3000	1500	35%	3153.87	1576.93	39.9%	41.9%
2	8	3000	1500	40%	3173.04	1586.52	46.4%	49.1%
2	9	3000	1500	45%	3191.65	1595.83	53.2%	56.8%
2	10	3000	1500	50%	3209.74	1604.87	60.2%	64.8%
3	1	3000	1500	5%	3036.82	1518.41	5.1%	5.1%
3	2	3000	1500	10%	3072.34	1536.17	10.4%	10.5%
3	3	3000	1500	15%	3106.67	1553.34	15.9%	16.2%
3	4	3000	1500	20%	3139.91	1569.95	21.6%	22.1%
3	5	3000	1500	25%	3172.11	1586.06	27.4%	28.4%
3	6	3000	1500	30%	3203.37	1601.68	33.5%	35.0%
3	7	3000	1500	35%	3233.74	1616.87	39.9%	41.9%
3	8	3000	1500	40%	3263.27	1631.64	46.4%	49.1%

3	9	3000	1500	45%	3292.03	1646.01	53.2%	56.8%
3	10	3000	1500	50%	3320.05	1660.02	60.2%	64.8%
4	1	3000	1500	5%	3049.19	1524.59	5.1%	5.1%
4	2	3000	1500	10%	3096.84	1548.42	10.4%	10.5%
4	3	3000	1500	15%	3143.07	1571.53	15.9%	16.2%
4	4	3000	1500	20%	3187.98	1593.99	21.6%	22.1%
4	5	3000	1500	25%	3231.65	1615.83	27.4%	28.4%
4	6	3000	1500	30%	3274.18	1637.09	33.5%	35.0%
4	7	3000	1500	35%	3315.63	1657.81	39.9%	41.9%
4	8	3000	1500	40%	3356.07	1678.03	46.4%	49.1%
4	9	3000	1500	45%	3395.55	1697.78	53.2%	56.8%
4	10	3000	1500	50%	3434.14	1717.07	60.2%	64.8%
5	1	3000	1500	5%	3061.61	1530.81	5.1%	5.1%
5	2	3000	1500	10%	3121.53	1560.77	10.4%	10.5%
5	3	3000	1500	15%	3179.89	1589.94	15.9%	16.2%
5	4	3000	1500	20%	3236.78	1618.39	21.6%	22.1%
5	5	3000	1500	25%	3292.31	1646.15	27.4%	28.4%
5	6	3000	1500	30%	3346.55	1673.28	33.5%	35.0%
5	7	3000	1500	35%	3399.59	1699.80	39.9%	41.9%
5	8	3000	1500	40%	3451.50	1725.75	46.4%	49.1%
5	9	3000	1500	45%	3502.34	1751.17	53.2%	56.8%
5	10	3000	1500	50%	3552.16	1776.08	60.2%	64.8%
6	1	3000	1500	5%	3074.09	1537.04	5.1%	5.1%
6	2	3000	1500	10%	3146.43	1573.21	10.4%	10.5%
6	3	3000	1500	15%	3217.14	1608.57	15.9%	16.2%
6	4	3000	1500	20%	3286.34	1643.17	21.6%	22.1%
6	5	3000	1500	25%	3354.10	1677.05	27.4%	28.4%
6	6	3000	1500	30%	3420.53	1710.26	33.5%	35.0%
6	7	3000	1500	35%	3485.69	1742.84	39.9%	41.9%
6	8	3000	1500	40%	3549.65	1774.82	46.4%	49.1%
6	9	3000	1500	45%	3612.48	1806.24	53.2%	56.8%
6	10	3000	1500	50%	3674.23	1837.12	60.2%	64.8%
7	1	3000	1500	5%	3086.61	1543.30	5.1%	5.1%
7	2	3000	1500	10%	3171.52	1585.76	10.4%	10.5%
7	3	3000	1500	15%	3254.83	1627.42	15.9%	16.2%
7	4	3000	1500	20%	3336.65	1668.32	21.6%	22.1%
7	5	3000	1500	25%	3417.06	1708.53	27.4%	28.4%
7	6	3000	1500	30%	3496.14	1748.07	33.5%	35.0%
7	7	3000	1500	35%	3573.96	1786.98	39.9%	41.9%
7	8	3000	1500	40%	3650.59	1825.29	46.4%	49.1%
7	9	3000	1500	45%	3726.08	1863.04	53.2%	56.8%
7	10	3000	1500	50%	3800.50	1900.25	60.2%	64.8%
8	1	3000	1500	5%	3099.18	1549.59	5.1%	5.1%
8	2	3000	1500	10%	3196.81	1598.40	10.4%	10.5%

8	3	3000	1500	15%	3292.96	1646.48	15.9%	16.2%
8	4	3000	1500	20%	3387.73	1693.86	21.6%	22.1%
8	5	3000	1500	25%	3481.19	1740.60	27.4%	28.4%
8	6	3000	1500	30%	3573.42	1786.71	33.5%	35.0%
8	7	3000	1500	35%	3664.46	1832.23	39.9%	41.9%
8	8	3000	1500	40%	3754.39	1877.20	46.4%	49.1%
8	9	3000	1500	45%	3843.26	1921.63	53.2%	56.8%
8	10	3000	1500	50%	3931.11	1965.56	60.2%	64.8%
9	1	3000	1500	5%	3111.81	1555.91	5.1%	5.1%
9	2	3000	1500	10%	3222.30	1611.15	10.4%	10.5%
9	3	3000	1500	15%	3331.54	1665.77	15.9%	16.2%
9	4	3000	1500	20%	3439.59	1719.80	21.6%	22.1%
9	5	3000	1500	25%	3546.53	1773.27	27.4%	28.4%
9	6	3000	1500	30%	3652.40	1826.20	33.5%	35.0%
9	7	3000	1500	35%	3757.26	1878.63	39.9%	41.9%
9	8	3000	1500	40%	3861.16	1930.58	46.4%	49.1%
9	9	3000	1500	45%	3964.12	1982.06	53.2%	56.8%
9	10	3000	1500	50%	4066.21	2033.10	60.2%	64.8%
11	1	3000	1500	5%	3137.22	1568.61	5.1%	5.1%
11	2	3000	1500	10%	3273.89	1636.95	10.4%	10.5%
11	3	3000	1500	15%	3410.05	1705.03	15.9%	16.2%
11	4	3000	1500	20%	3545.72	1772.86	21.6%	22.1%
11	5	3000	1500	25%	3680.91	1840.46	27.4%	28.4%
11	6	3000	1500	30%	3815.66	1907.83	33.5%	35.0%
11	7	3000	1500	35%	3949.97	1974.99	39.9%	41.9%
11	8	3000	1500	40%	4083.87	2041.94	46.4%	49.1%
11	9	3000	1500	45%	4217.37	2108.69	53.2%	56.8%
11	10	3000	1500	50%	4350.49	2175.25	60.2%	64.8%
12	1	3000	1500	5%	3150	1575	5.1%	5.1%
12	2	3000	1500	10%	3300	1650	10.4%	10.5%
12	3	3000	1500	15%	3450	1725	15.9%	16.2%
12	4	3000	1500	20%	3600	1800	21.6%	22.1%
12	5	3000	1500	25%	3750	1875	27.4%	28.4%
12	6	3000	1500	30%	3900	1950	33.5%	35.0%
12	7	3000	1500	35%	4050	2025	39.9%	41.9%
12	8	3000	1500	40%	4200	2100	46.4%	49.1%
12	9	3000	1500	45%	4350	2175	53.2%	56.8%
12	10	3000	1500	50%	4500	2250	60.2%	64.8%

Appendix B: Additional Data Analyses (WORKING PAPER)

B.1 Re-Estimation of Coller and Williams [1999]

Coller and Williams [1999] is a classic study, in the sense that it implemented many of the procedures that have become standard in the later literature. For that reason, it is useful to re-estimate their model, using their data. The reason is that we want to be clear what their data shows, and it is easy to draw ambiguous inferences from their statement of results. This is because many of the samples in specific treatments were small and involved significant differences in the demographic composition of the sample. For this reason one wants to control for demographics, simply to ensure that the marginal effects of treatments is clear.

Another reason for undertaking a re-estimation is that the econometric methods they employed, "interval regression" in which the dependent variable is recognized as only coming from an interval, were relatively new at the time. The econometric specification they used was "hand written" in LIMDEP, and is now standard in many econometric packages such as *Stata*. The popularity of this procedure means that newer statistical packages will likely generate better estimates (in large part from the extra work, "under the numerical hood," involved in finding good "starting values.") Related to this point, their econometric model used a multiplicative heteroskedasticity specification, to allow the residual variance to depend on covariates. As valuable as this extension is, it is known to generally lead to relatively flat likelihood functions, also demanding attention to numerical accuracy.

All variables are defined as in Coller and Williams [1999; p.119], apart from our variable `infomkt` which is the same as their cryptic `armkt`. Apologies for all output being in *Stata*-format, but this is an appendix marked "not for publication."

The first set of results use the interval regression model to conveniently summarize the data and then show total effects of each treatment. So we see from the first two sets of estimates, for example, that the overall discount rate, in annualized terms, is 36.5%, but that is it 72.5% when there is no front end delay and 42.3% lower when there is a front end delay (all percentages stated in this text should be understood to be percentage points).

```
. * total effects
Interval regression                               Number of obs   =        199
                                                  LR chi2(0)      =       0.00
Log likelihood = -789.33316                       Prob > chi2     =          .

------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [90% Conf. Interval]
-------------+----------------------------------------------------------------
       _cons |   36.51779     3.5657    10.24   0.000     30.65274    42.38285
-------------+----------------------------------------------------------------
    /lnsigma |   3.913105   .0539376    72.55   0.000     3.824385    4.001824
-------------+----------------------------------------------------------------
       sigma |   50.05412   2.699798                      45.80464    54.69785
------------------------------------------------------------------------------

Interval regression                               Number of obs   =        199
                                                  Wald chi2(1)    =       9.04
Log likelihood = -771.75144                       Prob > chi2     =     0.0026

------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [90% Conf. Interval]
-------------+----------------------------------------------------------------
model        |
```

```
         fed |  -42.29732   14.06864    -3.01   0.003    -65.43818   -19.15645
       _cons |   72.52874   13.68678     5.30   0.000     50.01599     95.0415
-------------+----------------------------------------------------------------
lnsigma      |
         fed |  -.5814246   .1481421    -3.92   0.000    -.8250966   -.3377525
       _cons |   4.320643   .1364622    31.66   0.000     4.096183    4.545103
------------------------------------------------------------------------------
```

```
Interval regression                               Number of obs   =        199
                                                  Wald chi2(1)    =       8.08
Log likelihood = -777.60117                       Prob > chi2     =     0.0045

------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [90% Conf. Interval]
-------------+----------------------------------------------------------------
model        |
       araer |  -20.54823   7.228204    -2.84   0.004    -32.43757   -8.658893
       _cons |   47.44993   6.145224     7.72   0.000     37.34194    57.55793
-------------+----------------------------------------------------------------
lnsigma      |
       araer |  -.4286368   .1074892    -3.99   0.000    -.6054409   -.2518327
       _cons |   4.084631   .0792418    51.55   0.000      3.95429    4.214972
------------------------------------------------------------------------------
```

```
Interval regression                               Number of obs   =        199
                                                  Wald chi2(1)    =       2.96
Log likelihood = -787.36913                       Prob > chi2     =     0.0853

------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [90% Conf. Interval]
-------------+----------------------------------------------------------------
model        |
        real |    14.1634    8.23214     1.72   0.085      .622739    27.70407
       _cons |   24.85965   7.181675     3.46   0.001     13.04685    36.67246
-------------+----------------------------------------------------------------
lnsigma      |
        real |   .1904092   .1397528     1.36   0.173    -.0394638    .4202821
       _cons |   3.746409   .1264302    29.63   0.000      3.53845    3.954369
------------------------------------------------------------------------------
```

The next set of estimates includes dummies for each of the treatments, but still pools across all subjects. We see that the front end delay has a significant effect of lowering elicited discount rates by 31 percentage points, and that this estimate has a *p*-value of 0.066. None of the other treatments have a statistically significant effect here, which is mildly disturbing until one corrects for demographics.

```
Interval regression                               Number of obs   =        199
                                                  Wald chi2(5)    =      12.41
Log likelihood = -762.28984                       Prob > chi2     =     0.0296

------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [90% Conf. Interval]
-------------+----------------------------------------------------------------
model        |
       araer |  -12.05499   11.73735    -1.03   0.304    -31.36122     7.25124
         mkt |  -12.99401   11.35876    -1.14   0.253     -31.6775    5.689485
        real |   .8000898   8.848733     0.09   0.928    -13.75478    15.35496
     infomkt |   8.916011   14.20957     0.63   0.530    -14.45666    32.28868
         fed |  -30.73606    16.7011    -1.84   0.066    -58.20692   -3.265201
       _cons |    71.7287   16.29811     4.40   0.000      44.9207     98.5367
-------------+----------------------------------------------------------------
lnsigma      |
       araer |  -.2849961   .1770697    -1.61   0.108    -.5762499    .0062576
         mkt |  -.5377438   .1889276    -2.85   0.004    -.8485021   -.2269856
        real |  -.4061529   .1848594    -2.20   0.028    -.7102195   -.1020864
     infomkt |   .1332218   .2611174     0.51   0.610    -.2962781    .5627218
         fed |  -.2908686   .1888792    -1.54   0.124    -.6015472    .0198099
       _cons |   4.726796   .2297714    20.57   0.000     4.348856    5.104737
------------------------------------------------------------------------------
```

The final estimation includes dummies and controls for individual demographic characteristics of the sample. The reason that this makes such a big difference here is that there were large differences in the sample composition of many of the sessions conducted here. In part this is due to some of the sessions have a relatively small sample size, and in part it is just due to the vagaries of recruitment. Of course, one might pursue these demographic differences further using a sample selection model, but that is beyond the scope (and data availability) for this re-estimation. At the very least one should simply correct for the differences in demographics, so that the marginal effect of the treatment is clearer.

These results, then, do lift a cloud of imprecision from two of the treatments that one might have expected to be significant. One is the provision of information on the implied annualized interest rates from each choice (variable `araer`), which now is shown to result in discount rates that are 12.9 percentage points lower (*p*-value of 0.050). The other is the use of real rewards rather than hypothetical survey questions (variable `real`), which raises elicited discount rates by 7.7 percentage points (*p*-value of 0.037). And, of course, the effect of the front end delay remains significant and large.

It is also worth noting that almost all of the treatments have a statistically significant effect on the residual variance of the dependant variable.

The demographics included here are the ones documented by Coller and Williams [1999; p.119], with one slight exception. There were 22 subjects that did not complete the question on "parental income" for one reason or another: hence the estimates in Coller and Williams [1999; Table 5, p. 120] only use 177 = 199 - 22 observations. We imputed parental income for these subjects at the median response from the 177 that did respond, and then formed a categorical variable so that the precise value of parental income was not assumed. One could use more elaborate methods, such as multiple imputation, for this step, but that seems overkill for this purpose.

```
Interval regression                                 Number of obs   =        199
                                                    Wald chi2(19)   =     109.16
Log likelihood = -679.95948                         Prob > chi2     =     0.0000

------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [90% Conf. Interval]
-------------+----------------------------------------------------------------
model        |
         age |   1.239518   .7854096     1.58   0.115    -.0523658    2.531402
        Male |   7.104897   2.819766     2.52   0.012     2.466796      11.743
        Race |   6.497151   3.891669     1.67   0.095     .0959247    12.89838
       CWhhy |    .189742   .1369622     1.39   0.166    -.0355408    .4150247
       PARy2 |   1.298354   11.39606     0.11   0.909     -17.4465    20.04321
       PARy3 |  -3.394432   9.591225    -0.35   0.723    -19.17059    12.38173
       PARy4 |  -9.085107   9.417319    -0.96   0.335    -24.57522    6.405003
       PARy5 |  -3.332124   9.320018    -0.36   0.721    -18.66219    11.99794
       PARy6 |  -3.094702   10.78107    -0.29   0.774    -20.82799    14.63859
       PARy7 |  -5.670802   9.822198    -0.58   0.564    -21.82688    10.48528
       PARy8 |   5.848071   10.26003     0.57   0.569    -11.02817    22.72432
       PARy9 |   6.212761   12.18999     0.51   0.610    -13.83799    26.26351
          hh |   12.87018   7.668459     1.68   0.093     .2566849    25.48367
         hh2 |  -2.686812   1.397189    -1.92   0.054    -4.984983   -.3886403
       araer |  -12.91904   6.577922    -1.96   0.050    -23.73876   -2.099323
         mkt |  -9.324936    6.63819    -1.40   0.160    -20.24379    1.593916
        real |   7.771713   3.730379     2.08   0.037     1.635785    13.90764
     infomkt |    9.23046   7.750076     1.19   0.234    -3.517281     21.9782
         fed |  -30.49949   14.66329    -2.08   0.038    -54.61845   -6.380528
       _cons |   14.00403   22.95289     0.61   0.542    -23.75012    51.75818
-------------+----------------------------------------------------------------
lnsigma      |
```

age	.0375082	.0198518	1.89	0.059	.0048548	.0701616
Male	.537626	.145972	3.68	0.000	.2975234	.7777285
Race	.7497739	.1780457	4.21	0.000	.4569148	1.042633
CWhhy	.0113375	.0041345	2.74	0.006	.0045369	.0181381
PARy2	-.2863221	.3613417	-0.79	0.428	-.8806763	.3080321
PARy3	-.7024058	.311774	-2.25	0.024	-1.215228	-.1895832
PARy4	-1.498469	.2983019	-5.02	0.000	-1.989132	-1.007806
PARy5	-1.197474	.2923123	-4.10	0.000	-1.678285	-.7166636
PARy6	-.6610539	.3271795	-2.02	0.043	-1.199216	-.1228916
PARy7	-.6085153	.3217474	-1.89	0.059	-1.137743	-.0792879
PARy8	-.1699196	.3015514	-0.56	0.573	-.6659276	.3260883
PARy9	-.0362514	.3067126	-0.12	0.906	-.5407488	.4682459
hh	1.557991	.5041435	3.09	0.002	.7287483	2.387233
hh2	-.3579855	.1083859	-3.30	0.001	-.5362644	-.1797066
araer	-.8166207	.2502693	-3.26	0.001	-1.228277	-.4049643
mkt	-1.066501	.2466513	-4.32	0.000	-1.472206	-.660796
real	.4802592	.2987445	1.61	0.108	-.0111317	.9716502
infomkt	.9138699	.3894284	2.35	0.019	.2733171	1.554423
fed	-.5885626	.2408822	-2.44	0.015	-.9847785	-.1923466
_cons	1.912509	.8147032	2.35	0.019	.5724418	3.252577

B.2 Expo-Power Utility Function

As explained in the text, an attractive generalization of the CRRA utility function, is the Expo-Power (EP) utility function proposed by Saha [1993]. Following Holt and Laury [2002], the EP function is defined as

$$U(x) = [1-\exp(-\alpha x^{1-r})]/\alpha, \tag{1}$$

where α and r are parameters to be estimated. RRA is then $r + \alpha(1-r)y^{1-r}$, so RRA varies with income if $\alpha \neq 0$. This function nests CRRA (as $\alpha \to 0$) and CARA (as $r \to 0$). Although we cannot formally reject the hypothesis of decreasing RRA, since $\alpha<0$, the variation in RRA over the domain of prizes presented to our subjects was minor. Here are the ML estimates, just using the responses to the lottery choice task:

(Std. Err. adjusted for 413 clusters in userid)

	Coef.	Robust Std. Err.	z	P>|z|	[95% Conf.	Interval]
mu						
_cons	.1786807	.0109985	16.25	0.000	.157124	.2002374
r						
_cons	.824448	.015911	51.82	0.000	.793263	.8556331
alpha						
_cons	-.2927343	.0290346	-10.08	0.000	-.349641	-.2358276

The range of RRA is then calculated for different prize levels, along with the 95% confidence interval. Figure B1 plots more detailed estimates of RRA. We therefore conclude that it is a reasonable approximation to assume CRRA in our analysis.

y	Point Estimate	95% Confidence	Interval
100	.7901455	.7611384	.8191525
200	.7857068	.7561253	.8152882
300	.7828487	.7528473	.8128501

400	.7806938	.7503519	.8110357
500	.7789458	.7483134	.8095782
1000	.773058	.7413594	.8047565
2000	.7664083	.7333585	.799458
3000	.7621264	.7281325	.7961203
4000	.7588981	.7241571	.7936392
5000	.7562794	.7209112	.7916476

Figure B1: Is Relative Risk Aversion Constant?

Maximum likelihood estimates of expo-power utility function
Predicted RRA and 95% confidence interval

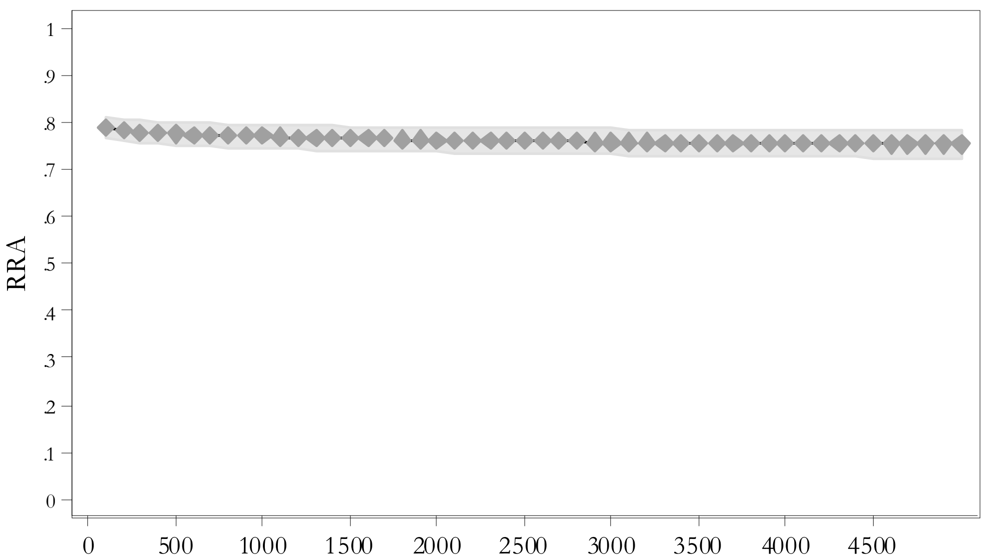

B.3 Analysis of Anderhub, Güth, Gneezy and Sonsino [2001]

Anderhub, Güth, Gneezy and Sonsino [2001] is a simple experiment in which discount rates can be inferred from the certainty equivalents of time-dated lotteries that are elicited from subjects. Assume that the Becker, DeGroot and Marschak [1964] procedure works reliably to elicit the certainty equivalent of a lottery.

In their P treatment they endowed each subject with 75 currency units (New Israeli Shekels), and asked them to state a *buying price* for a 50:50 lottery of 125 and 25 payable now, in 4 weeks, or in 8 weeks. These buying prices are certainty-equivalents for the lottery. Refer to the elicited certainty-equivalents for these time-dated lotteries as L0, L4 and L8. If these buying prices were below some randomly generated selling price, the subject kept the 75 now and did not get to play out the lottery. The ratio of L4 to L0 is a discount factor for a horizon of 4 weeks starting now, and the ratio of L8 to L4 is a discount factor for a horizon of 4 weeks starting in 4 weeks. From these revealed discount factors one can infer discount rates on an annualized basis, using the standard formulae. Then one can see if the discount rate between now and the 4 week horizon is different from the discount rate between 4 weeks and 8 weeks: in particular, given the one-sided prior that the literature has generated, one can test if these discount rates are declining. After making three valuation decisions, one of the 3 horizons was played out at random.

In their A treatment the subject was instead endowed with the lottery and asked to state a *selling price* for it. A computer-generated buying price was generated, and if the selling price was above that buying price the subject kept the lottery. Again, one of the three horizons was selected at random to be played out. And, again, the stated selling prices are the certainty-equivalents of the lottery.

The data for the P treatment is displayed below. Anderhub, Güth, Gneezy and Sonsino [2001; p.251-252] list the values of L0, L4 and L8, the ratio of L4 to L0, and the ratio of L8 to L4. In our notation p_L0 is the reported value of L0 in the P treatment, p_d1 is the reported ratio of L4 to L0 in the P treatment, and p_d2 is the reported ratio of L8 to L4. Hence, for subject #1, the value of L0 in the P treatment is 26, so the inferred values of L4 and L8 are also 26, since the ratio in each case is equal to 1. For subject #2, L4 is inferred to be 40 × 0.875 = 35, and L8 is inferred to be 35 × 0.857 = 30. We take the values of p_L0, p_d1 and p_d2 directly from the published data (hence there may be some trivial rounding errors in recovering integer-valued certainty equivalents).

We then infer the monthly discount rates as p_dr1 and p_dr2, and then the annualized discount rates, assuming monthly compounding, as P_DR1 and P_DR2. Finally, we define the difference in implied annual discount rates as P_DR_DIFF = P_DR2 - P_DR1. The null hypothesis from the Exponential discounting model is that P_DR_DIFF is zero, and the null hypothesis from the Hyperbolic discounting model is that P_DR_DIFF < 0.

id	p_L0	p_d1	p_d2	p_dr1	p_dr2	P_DR1	P_DR2	P_DR_DIFF
1	26	1.000	1.000	0.000	0.000	0.00	0.00	0.00
2	40	0.875	0.857	0.143	0.167	3.96	5.37	1.41
3	45	0.778	0.857	0.285	0.167	19.34	5.37	-13.96
4	45	0.888	0.857	0.126	0.167	3.16	5.37	2.21
5	45	0.888	1.000	0.126	0.000	3.16	0.00	-3.16
6	45	0.888	1.000	0.126	0.000	3.16	0.00	-3.16
7	50	0.800	0.857	0.250	0.167	13.55	5.37	-8.18
8	50	0.800	1.000	0.250	0.000	13.55	0.00	-13.55
9	50	0.900	0.889	0.111	0.125	2.54	3.10	0.56
10	50	0.900	0.889	0.111	0.125	2.54	3.10	0.56
11	50	0.900	0.889	0.111	0.125	2.54	3.10	0.56
12	50	0.900	0.889	0.111	0.125	2.54	3.10	0.56
13	50	0.940	0.957	0.064	0.045	1.10	0.69	-0.41
14	50	1.000	1.000	0.000	0.000	0.00	0.00	0.00
15	50	1.000	1.000	0.000	0.000	0.00	0.00	0.00
16	50	1.000	1.000	0.000	0.000	0.00	0.00	0.00
17	50	1.000	1.000	0.000	0.000	0.00	0.00	0.00
18	50	1.000	1.000	0.000	0.000	0.00	0.00	0.00
19	50	1.000	1.000	0.000	0.000	0.00	0.00	0.00
20	50	1.000	1.000	0.000	0.000	0.00	0.00	0.00
21	50	1.000	1.000	0.000	0.000	0.00	0.00	0.00
22	50	1.000	1.000	0.000	0.000	0.00	0.00	0.00
23	50	1.000	1.000	0.000	0.000	0.00	0.00	0.00
24	60	0.833	1.000	0.200	0.000	7.96	0.00	-7.96
25	65	1.000	1.000	0.000	0.000	0.00	0.00	0.00

26	70	1.000	1.000	0.000	0.000	0.00	0.00	0.00
27	75	1.000	1.000	0.000	0.000	0.00	0.00	0.00

There were 27 subjects in treatment P. The obvious thing to see from these data is that a large number of subjects exhibited no discount rate at all, and stated the same certainty-equivalent (e.g., #27). Whatever the procedural interpretation of these responses, and there could be many, we should evaluate the data with and without their responses.

The corresponding data from the 34 subjects in the A treatment is listed below, and is defined identically.

id	a_L0	a_d1	a_d2	a_dr1	a_dr2	A_DR1	A_DR2	A_DR_D~F
1	50	0.800	0.750	0.250	0.333	13.55	30.57	17.02
2	50	0.900	0.888	0.111	0.126	2.54	3.16	0.62
3	50	0.940	0.957	0.064	0.045	1.10	0.69	-0.41
4	60	0.666	0.875	0.502	0.143	130.31	3.96	-126.35
5	60	0.833	0.800	0.200	0.250	7.96	13.55	5.59
6	60	0.833	0.800	0.200	0.250	7.96	13.55	5.59
7	60	0.916	0.909	0.092	0.100	1.87	2.14	0.28
8	60	0.916	0.909	0.092	0.100	1.87	2.14	0.28
9	65	0.846	0.727	0.182	0.376	6.44	44.88	38.44
10	65	0.846	0.818	0.182	0.222	6.44	10.14	3.70
11	70	0.857	1.000	0.167	0.000	5.37	0.00	-5.37
12	70	0.957	0.970	0.045	0.031	0.69	0.44	-0.25
13	72	0.972	1.000	0.029	0.000	0.41	0.00	-0.41
14	75	0.933	0.971	0.072	0.030	1.30	0.42	-0.87
15	75	0.667	1.000	0.499	0.000	127.97	0.00	-127.97
16	75	0.800	0.667	0.250	0.499	13.55	127.97	114.42
17	75	0.800	0.833	0.250	0.200	13.55	7.96	-5.59
18	75	0.867	0.769	0.153	0.300	4.54	22.38	17.84
19	75	0.900	0.889	0.111	0.125	2.54	3.10	0.56
20	75	0.906	0.882	0.104	0.134	2.27	3.51	1.24
21	75	0.933	0.928	0.072	0.078	1.30	1.45	0.15
22	75	0.933	1.000	0.072	0.000	1.30	0.00	-1.30
23	75	1.000	1.000	0.000	0.000	0.00	0.00	0.00
24	75	1.000	1.000	0.000	0.000	0.00	0.00	0.00
25	75	1.000	1.000	0.000	0.000	0.00	0.00	0.00
26	75	1.000	1.000	0.000	0.000	0.00	0.00	0.00
27	75	1.000	1.000	0.000	0.000	0.00	0.00	0.00
28	75	1.000	1.000	0.000	0.000	0.00	0.00	0.00
29	75	1.000	1.000	0.000	0.000	0.00	0.00	0.00
30	75	1.000	1.000	0.000	0.000	0.00	0.00	0.00
31	75	1.000	1.000	0.000	0.000	0.00	0.00	0.00
32	75	1.066	1.062	-0.062	-0.058	-0.54	-0.51	0.02
33	100	0.800	0.875	0.250	0.143	13.55	3.96	-9.59
34	100	0.950	0.947	0.053	0.056	0.85	0.92	0.07

Again, we see a significant number of subjects that did not display any discount rate (e.g., #23), and indeed one subject (#32) that displayed a negative discount rate.

Overall discount rates are very, very high. Including the complete sample, we have the following statistics

```
    Variable |       Obs        Mean    Std. Dev.       Min        Max
-------------+--------------------------------------------------------
       P_DR1 |        27    2.929826    4.982409          0   19.33503
       P_DR2 |        27    1.281289     2.05776          0   5.371331

       A_DR1 |        34    10.84406    30.34475  -.5355771   130.3135
       A_DR2 |        34     8.71797    23.25862  -.5141462   127.9704
```

where these are rates, not percentage points. So the average annualized discount rate for the 4 weeks from now is 293% in the P treatment, and 1,084% in the A treatment. If we drop out the subjects that displayed no positive discount rate at all, things are even worse:

```
    Variable |       Obs        Mean    Std. Dev.       Min        Max
-------------+--------------------------------------------------------
       P_DR1 |        13    6.085023    5.743413   1.101192   19.33503
       P_DR2 |         9    3.843867    1.637144   .6945588   5.371331

       A_DR1 |        24    15.38473    35.32696   .4060631   130.3135
       A_DR2 |        20    14.84626    29.04389   .4235383   127.9704
```

Simple one-sample *t*-tests can be used to evaluate these data, focusing now on the differences in implied discount rates. For the P treatment we have, using the entire sample:

```
. ttest A_DR_DIFF=0

One-sample t test
------------------------------------------------------------------------------
Variable |     Obs        Mean    Std. Err.   Std. Dev.   [95% Conf. Interval]
---------+--------------------------------------------------------------------
A_DR_DIFF       34    -2.12609    6.513188    37.97809   -15.37727    11.12509
------------------------------------------------------------------------------
    mean = mean(A_DR_DIFF)                                        t =  -0.3264
Ho: mean = 0                                     degrees of freedom =       33

    Ha: mean < 0                 Ha: mean != 0                 Ha: mean > 0
 Pr(T < t) = 0.3731         Pr(|T| > |t|) = 0.7462          Pr(T > t) = 0.6269
```

From the perspective of the Exponential model the two-sided alternative hypothesis in the middle is the appropriate one, and the *p*-value of 0.74 clearly shows that one cannot reject the null that the Exponential discounting model is correct. Of course, this may be biased because of the large number of subjects that "flat-lined" with their certainty-equivalents and had 0 discount rates; we evaluate this possibility in a moment. From the perspective of the Hyperbolic model the one-sided alternative hypothesis on the left is the appropriate one, and the *p*-value of 0.37 clearly shows again that one cannot reject the null that the Exponential discounting model is correct. If we drop the observations with no discount rate, the sample size drops significantly, but we obtain these test results:

```
. ttest P_DR_DIFF=0 if p_dr1>0 & p_dr2>0

One-sample t test
------------------------------------------------------------------------------
Variable |     Obs        Mean    Std. Err.   Std. Dev.   [95% Conf. Interval]
---------+--------------------------------------------------------------------
P_DR_DIFF        9    -1.85339    1.823059    5.469176   -6.057371     2.35059
------------------------------------------------------------------------------
    mean = mean(P_DR_DIFF)                                        t =  -1.0166
Ho: mean = 0                                     degrees of freedom =        8

    Ha: mean < 0                 Ha: mean != 0                 Ha: mean > 0
 Pr(T < t) = 0.1695         Pr(|T| > |t|) = 0.3391          Pr(T > t) = 0.8305
```

So we arrive at the same qualitative conclusions as when we use the complete sample.

Turning to the A treatment, the two comparable sets of tests are as follows:

```
. ttest A_DR_DIFF=0

One-sample t test
------------------------------------------------------------------------------
Variable |     Obs        Mean    Std. Err.   Std. Dev.   [95% Conf. Interval]
---------+--------------------------------------------------------------------
A_DR_DIFF       34    -2.12609    6.513188    37.97809   -15.37727    11.12509
------------------------------------------------------------------------------
    mean = mean(A_DR_DIFF)                                        t =  -0.3264
Ho: mean = 0                                     degrees of freedom =       33

    Ha: mean < 0                 Ha: mean != 0                 Ha: mean > 0
 Pr(T < t) = 0.3731         Pr(|T| > |t|) = 0.7462          Pr(T > t) = 0.6269


. ttest A_DR_DIFF=0 if a_dr1>0 & a_dr2>0

One-sample t test
------------------------------------------------------------------------------
Variable |     Obs        Mean    Std. Err.   Std. Dev.   [95% Conf. Interval]
---------+--------------------------------------------------------------------
A_DR_DIFF       20    3.136883     9.05347    40.48835   -15.81225    22.08601
------------------------------------------------------------------------------
    mean = mean(A_DR_DIFF)                                        t =   0.3465
Ho: mean = 0                                     degrees of freedom =       19

    Ha: mean < 0                 Ha: mean != 0                 Ha: mean > 0
 Pr(T < t) = 0.6336         Pr(|T| > |t|) = 0.7328          Pr(T > t) = 0.3664
```

Hence there is no evidence from the A treatment that contradicts the Exponential discounting model.

We therefore conclude that the evidence in Anderhub, Güth, Gneezy and Sonsino [2001] is consistent with the Exponential discounting model, in keeping with their own non-parametric analysis (Table 1, p.245).

Additional References

Becker, Gordon M.; DeGroot, Morris H., and Marschak, Jacob., "Measuring Utility By A Single-Response Sequential Method," *Behavioral Science*, 9, July 1964, 226-232.

Saha, Atanu, "Expo-Power Utility: A Flexible Form for Absolute and Relative Risk Aversion," *American Journal of Agricultural Economics*, 75(4), November 1993, 905-913.

Appendix C: Random Coefficients (WORKING PAPER)

Our random coefficients specification directly estimates the structural parameters of the utility functions and discounting functions reviewed in the text. The conventional random coefficients specification, also referred to as a "mixed specification," assumes a linear latent index. We need to generalize that specification to allow non-linear latent indices. Following Andersen, Harrison, Hole, Lau and Rutström [2012], we focus on the basic logic assuming we are just trying to estimate one coefficient in a simple set of risk aversion tasks. The logic extends immediately to the joint estimation setting that our analysis requires when we include the choices from the discounting tasks. Moreover, it extends immediately to allow mixture specifications.

C.1 Basic Random Coefficients Specification

Assume a sample of N subjects making choices over J lotteries in T experimental tasks.[35] In all of the applications we consider, J=2 since the subjects are making choices over two lotteries or time-dated payments, but there are many designs in which the subject is asked to make choices over J>2 lotteries (e.g., Binswanger [1981]). In the traditional mixed logit literature one can view the individual n as deriving random utility Δ from alternative j in task t, given by

$$\Delta_{njt} = \beta_n x_{njt} + \boldsymbol{\varepsilon}_{njt} \tag{1}$$

where β_n is a vector of coefficients specific to subject n, x_{njt} is a vector of observed attributes of individual j and/or alternative j in task t, and $\boldsymbol{\varepsilon}_{njt}$ is a random term that is assumed to be identically and independently distributed extreme value. We use the symbol Δ for utility in (1), since we will need to generalize to allow for non-linear utility and discounting functions, and prefer to think of (1) as defining a latent index rather than as utility. In our experience, this purely semantic difference avoids some confusions about interpretation.

Specifically, for our purposes we need to extend (1) to allow for non-linear functions G defined over β and the values of x, such as

$$\Delta_{njt} = G(\beta_n, x_{njt}) + \boldsymbol{\varepsilon}_{njt} \tag{2}$$

For example, x might consist of the vector of monetary prizes m_k and probabilities p_k, for outcome k of K in a given lottery, and we might assume a Constant Relative Risk Aversion (CRRA) utility function

$$U(m_k) = m_k^{\ r} \tag{3}$$

where r is a parameter to be estimated.[36] Under expected utility theory (EUT) the probabilities for each outcome are those that are induced by the experimenter, so expected utility is simply the probability weighted utility of each outcome in each lottery j:

$$EU_j = \Sigma_k \,[\, p_k \times U(m_k) \,] \tag{4}$$

---

[35] It is trivial to allow J and T to vary with the individual, but for ease of notation we omit that generality.

[36] The choice of the power function is purely for pedagogical reasons and to keep the exposition simple.

If we let β=r here, we will want to let $G(\beta_n, x_{njt})$ be defined as

$$G(r_n, m_{njt}, p_{njt}) = EU_j \tag{5}$$

using (3) and (4), and hence let the latent index Δ in (2) be evaluated.[37]

The population density for β is denoted $f(\beta | \theta)$, where θ is a vector defining what we refer to as the hyper-parameters of the distribution of β. Thus individual realizations of β, such as β_n, are distributed according to some density function f. For example, if f is a Normal density then θ_1 would be the mean of that density and θ_2 the standard deviation of that density, and we would estimate the hyper-parameters θ_1 and θ_2. Or f could be a Uniform density and θ_1 would be the lower bound and θ_2 would be the upper bound. If β consisted of more than two parameters, then θ might also include terms representing the covariance of those parameters.

Conditional on β_n, the probability that the subject n chooses alternative i in task t is then given by the conditional logit formula, modestly extended to allow our non-linear index

$$L_{nit}(\beta_n) = \exp\{G(\beta_n, x_{nit})\} \;/\; \Sigma_j \exp\{G(\beta_n, x_{njt})\} \tag{6}$$

The probability of the observed choices by subject n, over all tasks T, again conditional on knowing β_n, is given by

$$P_n(\beta_n) = \Pi_t \; L_{ni(n,t)t}(\beta_n) \tag{7}$$

where i(n,t) denotes the lottery chosen by subject n in task t, following the notation of Revelt and Train [1998]. The unconditional probability involves integrating over the distribution of β:

$$P_n(\theta) = \int P_n(\beta_n) f(\beta | \theta) \, d\beta \tag{8}$$

and is therefore the weighted average of a product of logit formulas evaluated at different values of β, with the weights given by the density f.

We can then define the log-likelihood by

$$LL(\theta) = \Sigma_n \ln P_n(\theta) \tag{9}$$

and approximate it numerically using simulation methods, since it cannot be solved analytically. Using the methods of Maximum Simulated Likelihood (MSL) reviewed in Train [2003; §6.6, ch.10] and Cameron and Trivedi [2005; ch.12], we define the simulated log-likelihood by taking h=1,...,H replications β^h from the density $f(\beta | \theta)$:

---

[37] This approach generalizes immediately to non-EUT models in which there are more parameters, say to account for probability weighting and loss aversion. It also generalizes to non-CRRA specifications within EUT models that allow for more flexible specifications of risk attitudes that might vary with the level of the prizes. Each of these extensions involves more non-linearities than our EUT example, taking us even further from the domain of linear mixed logit.

$$SLL(\theta) = \sum_n \ln \{ \sum_h P_n(\beta^h)/H \} \tag{10}$$

The core insight of MSL is to evaluate the likelihood conditional on a randomly drawn β^h, do that H times, and then simply take the unweighted average over all H likelihoods so evaluated. The average is unweighted since each replication h is equally likely, by design. If H is "large enough," then MSL converges, under modest assumptions, to the Maximum Likelihood (ML) estimator.[38]

The value of this extension to non-linear mixed logit might not be obvious, because of widespread reliance on theorems showing that the linear mixed logit specification can approximate arbitrarily well any random-utility model (McFadden and Train [2000]; Train [2003; §6.5]).[39] So, why does one need a non-linear mixed logit specification? The reason is that these results only go in one direction: for any specification of a latent structure, defined over "deep parameters" such as risk preferences, they show that there exists an equivalent linear mixed logit. But they do not allow the direct recovery of those deep parameters in the estimates from the linear mixed logit. The deep parameters, which may be the things of interest, are buried in the estimates from the mixed logit, but can only be identified with restrictive assumptions about functional form. For example, risk attitudes can be considered using a linear specification if one assumes that utility is quadratic or that the distribution of returns are Normal (e.g., Luenberger [1998; §9.5]); neither are palatable assumptions in general.

Our specification has been couched in the language of estimating the structural parameters of a model of risk attitudes, but is perfectly general. Another obvious example would be the use of technology or transportation choices to recover the structural parameters of production functions, or the use of stated or revealed choices to recover the structural parameters of utility functions defined over consumption goods. The analyst needs to fill in their own equations for our (3) and (4), but only

---

[38] An important practical consideration with MSL is the manner in which replicates are drawn, and the size of H that is practically needed. We employ Halton draws to provide better coverage of the density than typical uniform number generators: see Train [2003; ch.9] for an exposition, and Drukker and Gates [2006] for the numerical implementation we employ. All results use H=250, which is generally large in relation to the literature. Our computational implementation generalizes the linear mixed logit program developed for *Stata* by Hole [2007].

[39] There are reasons to be suspicious of these theorems, although that is not critical for the point being made here. Specifically, two critical assumptions seem to connect observables and unobservables in a highly restrictive way. In one case, the correct claim is made (p.449) that a "primitive postulate of preference theory is that tastes are established prior to assignment of resource allocations." But this does not justify the assumption that "consumers with similar observed characteristics will have similar distributions of unobserved characteristics." Then a related, second assumption is made about attributes. Here the correct claim is that another "primitive postulate of consumer theory is that the description of a resource allocation does not depend on consumer characteristics. Thus, consumers' tastes and perceptions do not enter the 'objective' description of a resource allocation, although they will obviously enter the consumer's evaluation of the allocation." But it does not follow from this observation that "discrete alternatives that are similar in their observed attributes will have similar distributions of unobserved attributes." These assumptions are akin to the identifying assumptions of "random effects" specifications, that the random effect is orthogonal to the observed characteristics used as regressors. One other concern with these theorems is that they rest on polynomial approximations to random utility (McFadden and Train [2000; p. 466]), and these are known to have unreliable properties in statistical applications (e.g., White [1980; §2]). Referring to the class of approximations, including the polynomial, that are generated by applications of Taylor's Theorem, Gallant [1981; p. 212] notes that this "... theorem fails rather miserably as a means of understanding the statistical behavior of parameter estimates and test statistics."

need in the end to define $G(\beta_n, x_{njt})$ in (5) and hence in (2).

C.2 Flexible Population Distributions "For Free"

"But there is more," as they say on those kitchy television commercials for the latest set of super-knives! In principle the mixed logit specification, whether linear or non-linear, allows a wide range of shapes for the probability distribution used to characterize the population. In practice, one typically sees a relatively simple set of distributions used: univariate or multivariate Normal distributions, log-Normal distributions for coefficients known to be non-negative, uniform distributions, or triangular distributions.

One attractive option, since we are already allowing non-linear transformations of the population parameters, is to employ a transformation of the Normal distribution known as the Logit-Normal (L-N) distribution. Originally proposed by Aitchison and Begg [1976; p.3] as an excellent, tractable approximation to the Beta distribution, it has been resurrected by Lesaffre, Rizopoulos and Tsonaka [2007]. One nice property of the L-N distribution is that MSL algorithms developed for univariate or multivariate Normal distributions can be applied directly, providing one allows non-linear transformations of the structural parameters, and that is exactly what we are doing already to estimate structural parameters.

Figures C1 and C2 illustrate the wide array of distributional forms that are accommodated by the L-N distribution. The bi-modal and skewed distributions that are possible are particularly attractive. Note that these alternatives are all generated by different values of the two parameters of the (univariate) Normal distribution, so there is no "extra cost" of this flexibility in terms of additional parameters.

One limitation, of course, is that the Beta distribution and the L-N approximation of it, are defined over the unit interval. For some important inferential purposes, such as estimating a subjective probability, this is not a concern, but in general we would like something that is more general. In many other cases though, one would want the estimated distribution to be constrained to lie within specific boundaries dictated by theory. Examples include non-negativity constraints to ensure monotonicity and non-satiation in utility, or restrictions to the unit interval for probabilities or shares. In fact, the power utility function (3) that we employ here for illustration requires that $r>0$ to ensure monotonicity. It is a simple matter to define the so-called "Beta4 distribution" with two additional parameters: one to stretch out the distribution or squeeze it up, and another parameter to shift it left or right. This flexibility makes it possible to theoretically constrain the distribution of the structural parameter to be estimated.[40]

---

[40] Obviously some constraints can be accommodated by well-known transformations, such as non-negativity and the Log-normal. This alternative is often standard in linear mixed logit specifications (e.g., Hole [2007; p.390] and the "ln" option). Our approach is more general, particularly for estimates constrained *a priori* to some finite interval.

C.3 Extensions for Discounting and Mixture Models

It is an immediate extension of the previous specification to extend the set of binary choices to include discounting choices, and allow for joint estimation of parameters. A particularly attractive feature of the random coefficients approach is that one can allow for correlation between population characteristics, and estimate it.

One extension which is conceptually immediate, but requires extended programming, is to allow for mixture specifications. The *Stata* programs developed by Andersen, Harrison, Hole, Lau and Rutström [2012] allow the user to specify a "utility function" using any specification in terms of the economics that is desired. This allows specification of a wide range of utility functions (e.g., CRRA, CARA, HARA, Expo-power), and indeed alternative decision-making models (e.g., RDU or PT). They also provided templates for users wanting to see how to code alternatives like these. But as a programming matter the "user interface" for this flexibility is "downstream" of the evaluation of the probability of choices condition on parameter values. That is, the user-friendly interface available for writing these alternative specifications does not actually evaluate the probability (and hence likelihood) of the trial value of parameters: this is undertaken "upstream" in some elegant code that is likely cryptic to outsiders. Given the interest in mixture specifications, we therefore extended the *Stata* programs to allow this in a user-friendly manner. The user defines two "utility functions" instead of one, and the final parameter in the program is assumed to be the mixing probability. Instead of making a call to the *Stata* command `mixlognl`, the call is to the *Stata* command `mixlognlmm`. The programming logic can be easily extended to mixtures of three or more models, although that is not needed for our analysis.[41]

Additional References

Aitchison, J., and Begg, C.B., "Statistical Diagnosis when Basic Cases are Not Classified with Certainty," *Biometrika*, 63(1), 1976, 1-12.

Andersen, Steffen; Harrison, Glenn W.; Hole, Arne Rise; Lau, Morten I., and Rutström, E. Elisabet, "Non-Linear Mixed Logit," *Theory and Decision*, 73, 2012, 77-96.

Binswanger, Hans P., "Attitudes Toward Risk: Theoretical Implications of an Experiment in Rural India," *Economic Journal*, 91, December 1981, 867-890.

Cameron, A. C., and Trevedi, K.P., *Microeconometrics: Methods and Applications* (New York: Cambridge University Press, 2005).

Drukker, David M., and Gates, Richard, "Generating Halton sequences using Mata," *Stata Journal*, 6(2), 2006, 214-228.

Gallant, A. R., "On the Bias in Flexible Functional Forms and an Essentially Unbiased Form," *Journal of Econometrics*, 15, 1981, 211-245.

[41] Anyone wanting the *Stata* code for three or more models should send a check for $9.99E+6 payable to CEAR, to CEAR, Georgia State University, P.O. BOX 4036, Atlanta, GA 30302-4036, USA.

Hole, Arne Risa, "Fitting Mixed Logit Models by Using Maximum Simulated Likelihood," *Stata Journal*, 7(3), 2007, 388-401.

Lesaffre, Emmanuel; Rizopoulos, Dimitris, and Tsonaka, Roula, "The Logistic Transform for Bounded Outcome Scores," *Biostatistics*, 8(1), 2007, 72-85.

Luenberger, David G., *Investment Science* (New York: Oxford University Press, 1998).

McFadden, Daniel, and Train, Kenneth, "Mixed MNL Models for Discrete Response," *Journal of Applied Econometrics*, 15, 2000, 447-470.

Revelt, David, and Train, Kenneth, "Mixed Logit With Repeated Choices: Households' Choices of Appliance Efficiency Levels," *Review of Economics & Statistics*, 80, 1998, 647-657.

Train, Kenneth E., *Discrete Choice Methods with Simulation* (New York: Cambridge University Press, 2003).

White, Halbert, "Using Least Squares to Approximate Unknown Regression Functions," *International Economic Review*, 1980, 21, 149-170.

Figure C1: Illustrative Symmetric Logit-Normal Distribution

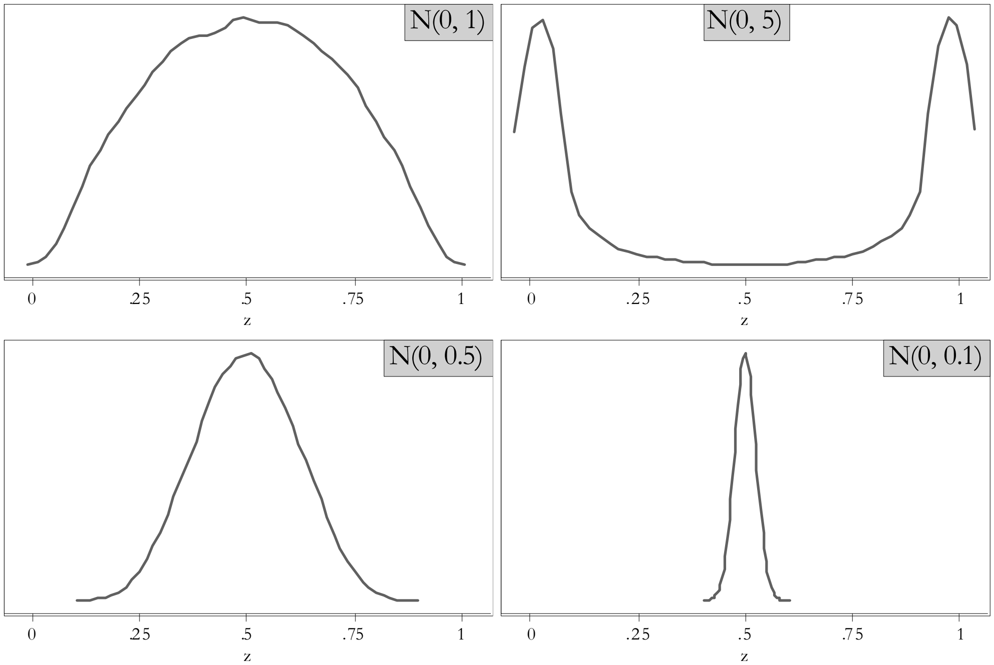

Figure C2: Illustrative Asymmetric Logit-Normal Distribution

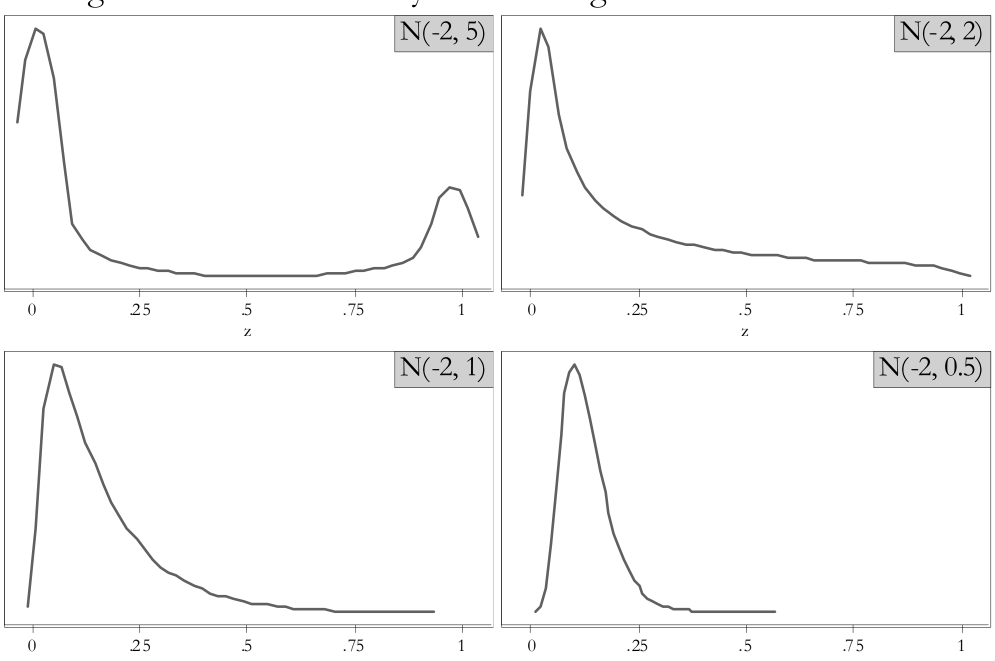

Appendix D: Literature Review (WORKING PAPER)

A. Discounting Models

The exposition closely follows §1 of the main text, but provides additional formal statements of discount rates implied by each model, rationales behind each model, and variants in the literature.

We define the discount factor for a given horizon τ to be the scalar D that equates the utility of the smaller level of income y received at time t with the larger level of income Y received at time $t+\tau$:

$$U(y_t) = D\, U(Y_{t+\tau}) \tag{0}$$

for some utility function U(.) and y < Y. We assume that the same utility function is used to evaluate income at time t and income at time $t+\tau$; we discuss this assumption later. This general definition of D permits the special case, much studied in the experimental literature, in which U(.) is linear. The non-linear case is of great empirical significance for inferences about discount rates, as demonstrated by Andersen, Harrison, Lau and Rutström (AHLR) [2008a]. There is nothing in (0) that restricts us to EUT, and indeed non-EUT specifications are considered later. We define utility over income and not directly over consumption flows or wealth, and discuss the implications of that specification later.

The discount factor for the Exponential (E) specification is defined as

$$D^E(t) = 1/(1+\delta)^t \tag{1}$$

for $t \geq 0$, and where the discount rate d is simply

$$d^E(t) = \delta \tag{2}$$

Although these characterizations are abstract, we view the discount rate on an annualized basis throughout. The key feature of this model, of course, is that the discount rate is a constant over time. The percentage rate at which utility today and utility tomorrow is discounted is exactly the same as the rate at which utility in 7 days and utility in 8 days is discounted. The debate over climate change has reminded us all that, with this specification, even small discount rates can lead to very low weight being placed on longer-term future consequences.

The discount factor for the Quasi-Hyperbolic (QH) specification is defined as

$$D^{QH}(t) = 1 \qquad \text{if } t = 0 \tag{3a}$$

$$D^{QH}(t) = \beta/(1+\delta)^t \qquad \text{if } t > 0 \tag{3b}$$

where $\beta<1$ implies quasi-hyperbolic discounting and $\beta=1$ is exponential discounting. Although the δ in (3b) may be estimated to be a different value than the δ in (1), or other specifications below, we use the same notation to allow comparability of functional forms. The defining characteristic of the QH specification is that the discount factor has a jump discontinuity at t=0, and that is thereafter exactly the same as the E specification. The discount rate for the QH specification is the value of $d^{QH}(t)$ that solves $D^{QH}(t) = 1/(1+d^{QH})^t$, so it is

$$d^{QH}(t) = [\ \beta/(1+\delta)^t\ ]^{(-1/t)} - 1 \tag{4}$$

for t>0. Thus for β<1 we observe a sharply declining discount rate in the very short run, and then the discount rate asymptotes towards δ as the effect of the initial drop in the discount factor diminishes. The drop β can be viewed as a fixed utility cost of discounting anything relative to the present, since it does not vary with the horizon t once t>0.[42] The QH specification was introduced by Phelps and Pollak [1968] for a social planning problem, and applied to model individual behavior by Elster [1979; p.71] and then Laibson [1997].

There are alternative ways to think of the fixed cost of discounting. Instead of thinking of the fixed cost as a percentage of the principal, one could think of it as a fixed monetary amount. The discount factor for the resulting Fixed Cost (FC) specification is defined as

$$D^{FC}(t) = 1 \qquad \text{if } t = 0 \tag{5a}$$

$$D^{FC}(t) = \beta\ [1 - (1-\theta)\delta t]^{\ (1/(1-\theta))} - (b/y_t) \qquad \text{if } t > 0 \tag{5b}$$

where β<1 indicates that there is a quasi-hyperbolic component to discounting, b>0 indicates that there is a fixed *monetary* cost component to discounting, and θ allows a wide range of discounting functions since θ=1 (with β=1 and b=0) implies exponential discounting, θ=2 (with β=1 and b=0) implies a form of hyperbolic discounting. The discount rate for the FC specification is

$$d^{FC}(t) = [\ \beta\ (1 - (1-\theta)\delta t)^{\ (1/(1-\theta))} - (b/y_t)\ ]^{\ (-1/t)} - 1 \tag{6}$$

for t>0. An obvious variant on (5b) is to allow the fixed cost component to be defined in terms of utility[43], since we are multiplying it by the utility of income, so we would have

$$D^{FC}(t) = \beta\ [1 - (1-\theta)\delta t]^{\ (1/(1-\theta))} - (b/U(y_t)) \qquad \text{if } t > 0 \tag{5c}$$

Of course, it is behaviorally possible that individuals behave as if they require some nominal amount of money before they delay receipt of income, implying that (5b) may be a better representation of behavior than (5c). The FC specification was proposed by Benhabib, Bisin and Schotter [2010].

There have been whole families of "hyperbolic" specifications of the discounting function. The simplest assumes a discount factor given by

$$D^{H1}(t) = 1/t \tag{7}$$

[42] One generalization of the QH specification is to allow there to be a jump discontinuity in the discount factor for some t=τ>0 rather than at t=0. Another is to allow the jump discontinuity to be ωβ instead of β, and for the exponential discounting term, that applies after time τ, to be weighted by (1-ω). The identification problems implied by the latter generalization are profound. The former generalization, by itself, could be evaluated by varying the front end delay continuously instead of the discrete variations we employed.

[43] One would then require that utility be unique up to positive affine transformations, rather than merely order-preserving transformations.

with discount rates $$d^{H1}(t) = t^{(1/t)} - 1 \tag{8}$$

The H1 specification was proposed in this manner by Ainslie [1975; p.472, Figure 3] as a direct translation of the "matching rule" specification of Herrnstein [1961; p.270] and Chung and Herrnstein [1967; p.70, equation (1)] for the delayed responses of animals to reward. This function has the obvious theoretical problem for sufficiently long horizons of allowing any finite change in utility in the near future to be offset by arbitrarily small changes of opposite sign in the future; of course, if the objective is to only make inferences over shorter horizons then this unfortunate property is just a *curiosem*. This theoretical problem is overcome by a simple generalization by Harvey [1986] discussed below.

A slight generalization of (7) is given by

$$D^{H2}(t) = 1/(1+ K\ t) \tag{9}$$

for some parameter K, with discount rates

$$d^{H2}(t) = (1 + K\ t)^{(1/t)} - 1, \tag{10}$$

and a further generalization by

$$D^{H3}(t) = 1/(1+ K\ t^{\varkappa}) \tag{11}$$

for some additional parameter $\varkappa$, and discount rates

$$d^{H3}(t) = (1 + K\ t^{\varkappa})^{(1/t)} - 1, \tag{12}$$

The H2 specification was first proposed by Mazur [1984; p.427], and the H3 specification by Mazur [1987; p.59].[44] An alternative generalization of (7) is

$$D^{H4}(t) = (1/t)^{r} \tag{13}$$

where $r>0$ is a parameter that determines the importance of the future, and with discount rates

$$d^{H4}(t) = (1/t)^{(-r/t)} - 1 \tag{14}$$

The primary attraction of (13) is that it implies much slower discounting of long-run consequences than the E specification (1), and avoids the awkward theoretical implications of (1) for longer-run planning. The H4 specification was proposed by Harvey [1986; p.1124, equation (2)], who differentiated it from (1) on an axiomatic basis. Essentially, (1) follows when time preferences are

---

[44] Mazur [1987; p.59] credits the idea of an exponent on the reinforcement stimuli as being due to a much older literature in psychology. Some of that literature views the psychological process underlying the parameter $\varkappa$ as simply reflecting the complexities of classical conditioning responses, and other parts of the literature see it as reflecting a more nuanced operant conditioning response by the subject to the stimuli. The latter interpretation anticipates the interpretation of the Weibull discounting function below. In general, Mazur [1987; p.72] was honestly agnostic about the psychological interpretation underlying the parameters in (9) and (11).

defined over the proportional changes in utility in two distinct time intervals of *equal length*, whereas (7) and (13) follow when time preferences are defined over the proportional changes in utility in two distinct time intervals of *proportionate length*. For example, (1) posits the decision maker comparing the percentage reduction in utility between years 5 and 6 and equating it to the percentage reduction in utility between years t and t+τ for τ=1 and ∀t. But (7) and (13) posit the decision maker comparing the percentage reduction in utility between years 5 and 10 and equating it to the percentage reduction in utility between years t and t+τ for (t+τ)/t = (10/5) = 2 and ∀t. These alternatives can be usefully viewed as different behavioral assumptions about how individuals cognitively compare utility streams across periods. Specification (13) then extends (7) to allow for different weight to be given to the future.[45]

One hyperbolic generalization of (7) is a variant of H2 and H4:

$$D^{H5}(t) = [1/(1+\alpha t)^{(\beta/\alpha)}] \tag{15}$$

for α, β > 0, and with discount rates

$$d^{H5}(t) = (1 + \alpha t)^{(\beta/\alpha t)} - 1 \tag{16}$$

This specification nests the E specification as $\alpha \rightarrow 0$, and was proposed by Loewenstein and Prelec [1992; p. 580]. Additional hyperbolic generalizations are proposed by Jamison and Jamison [2011].

A flexible specification is based on the Weibull (W) distribution from statistics[46], and is defined as

$$D^{W}(t) = \exp(-\acute{r}t^{(1/\acute{s})}) \tag{17}$$

for ŕ>0 and ś>0. For ś=1 this collapses to the E specification, and hence the parameter ś can be viewed as reflecting the "slowing down" or "speeding up" of time as perceived by the individual. This specification is due to Read [2001; p.25, equation (16)], although he noted (p.25, equation (15)) that the same point about time perception was implicit in the earlier hyperbolic generalization (10).[47] The discount rate at time t in this specification is then

$$d^{W}(t) = \exp(\acute{r}t^{(1-\acute{s})/\acute{s}}) - 1 \tag{18}$$

A further generalization of (17) is to think of a more general function of time as capturing the individual's perception of time, such as

---

[45] Harvey [1986; p. 1130, equation (7′)] anticipates the simple extension needed to allow for a positive front end delay. Harvey [1994; p.34, equation (2)] proposes a "proportional discounting" model in which the discount factor is b/(b+t) for some parameter b>0; the implied discount rates are $[b/(b+t)]^{-1/t} - 1$. However, he explicitly warns (p. 35) that this "... is a prescriptive model that is sufficiently simple to be applied in public studies. It is not intended as a descriptive model."

[46] Any probability density function f(t) defined on [0, ∞) can form the basis of a discounting function by taking the integral of f(t) between t and ∞.

[47] The W specification is the same as the simple functional form proposed in Prelec [2004; p. 526] and applied in Ebert and Prelec [2007; p. 1424ff.] and AHLR [2008a; p. 607].

$$D^W(t) = \exp(-ŕ^{c \ln(t)}) \quad (19)$$

for some parameter c. This specification was proposed by Roelofsma [1996; p.14, equation (3)] based on the psychometric function known as Weber's Law that the perceived difference between two sensory stimuli (points in time) is some constant proportion of the absolute magnitude of the stimuli.[48]

More generally, the literature on the Weibull discounting function suggests a parallel to the literature on probability weighting and decision weights in models of decision making under a-temporal risk. In the latter case the extension has been to allow decision-makers to treat objective probabilities as subjectively perceived, as well as to allow decision-makers to treat objective outcomes and payoffs as subjectively perceived through a utility function. This naturally leads to questions about how much of any observed risk aversion can be attributed to each component. In the discounting context, the psychology literature similarly talks about explanations of discounting behavior as being either "perceived-value-based accounts" or "perceived-time-based accounts" (Kim and Zauberman [2009; p. 92]. The former has often been interpreted as referring only to the δ parameter in an Exponential discounting specification defined over flows of money, and not defined over flows of utility. The latter has led to a small cottage industry of suggestions for functional forms, akin to the cottage industry experienced for probability weighting functional forms.[49]

For all of the formal specifications, there are some major themes that differentiate discounting models. For our purposes we want to focus on the exemplars of each approach, to avoid distraction with the specifics of each formulation. Obviously the E model (1) should be included as a benchmark, and the QH model (3a)-(3b) because of its popularity in behavioral economics. For the same reason, the FC model (5a)-(5b) should be considered. Within the family of "smooth" non-constant discounting models, the W specification (19) is attractive and flexible.

B. Older Experimental Studies

Ainslie and Haendel [1983; p.131-133] report experiments with 18 patients in a substance abuse program that made 66 choices over several weeks. Each subject earned a certain amount of money in an unrelated task during the week, ranging from $2 up to $10; call that $x. They were given a choice between receiving that $x in 7 days, or receiving $1.25x in 10 days. Then, on the 7th day, they were given a choice between receiving $x on that day or receiving $1.25x in 3 days. It is implied that this second choice was for the same $x, and not an additional choice, so subjects were allowed to change their minds on the 7th day. Observed behavior was generally consistent with exponential discounting for the majority of choices: 35% of the choices were consistently for the earlier option, and 27% of the choices were consistently for the later option. On the other hand, one-third were

[48] A similar approach is employed by Bleichrodt, Rohde and Wakker [2009; p. 31] who propose the Constant Absolute Decreasing Impatience (CADI) function $D^{CADI}(t) = \exp(ŕ (\exp (-ct)))$ for $c>0$; $D^{CADI}(t) = \exp(-ŕt)$ for $c=0$; and $D^{CADI}(t) = \exp(-ŕ (\exp (-ct)))$ for $c<0$, and some parameters $ŕ, c > 0$. A comparable specification known as the Constant Relative Decreasing Impatience (CRDI) is specified by Bleichrodt, Rohde and Wakker [2009; p. 32] as $D^{CRDI}(t) = \exp(ŕt^{1-d})$ for $d>1$; $D^{CRDI}(t) = t^{-ŕ}$ for $d=1$; and $D^{CRDI}(t) = \exp(-ŕ\, t^{1-d})$ for $d<1$.

[49] For example, Takahashi [2005] proposed the logarithmic function, Killeen [2009] used the power function, and Zauberman, Kim, Malkoc and Bettman [2009] used both. Scholten and Read [2006; Table 1] propose a "discounting by intervals" specification which generalizes (1), (11), (15) and (19).

consistent with hyperbolic or quasi-hyperbolic preferences, and entailed a shift from the later option to the sooner option. Thus there is, overall, evidence in favor of non-constant discounting, but for a minority of the observed choices.[50] It is not possible to draw any inferences about average discount rates from these data.

Horowitz [1991] was a remarkable early study that elicited a willingness to pay using a multiple-unit analogue of the Vickrey auction where the winning bidders pay the highest rejected bid.[51] The object was a fixed $50 to be paid in 64 days in one experiment, and in 34 days in an experiment conducted a month after the first experiment. If the weakly dominant strategy is understood, and this is behaviorally problematic in this sealed-bid context (Rutström [1998] and Harstad [2000]), the bid is the individual's true certainty-equivalent for the delayed payment. Hence the bid can be directly used to infer a discount factor and hence an annualized discount rate. In effect, then, the elicitation of the certainty-equivalent with this "uniform-price" sealed-bid auction bypasses the need for correction for non-linear utility, but at the price that the mechanism is notoriously hard to get to work behaviorally. The average (median) discount rate for the shorter horizon was 830% (436%) and for the longer horizon it was 271% (167%). Horowitz [1991; p. 320] is remarkably honest about the credibility and comprehension problems in his procedures, so these results should be taken with a pound of salt, but they clearly exhibit rampant, hyperbolicky behavior.

Kirby and Maraković [1996] undertook a clean experiment in which each subject was asked to make 21 binary choices between a certain amount of money today and a larger amount of money in the future. The "principal" was varied from choice to choice, as they were presented to subjects, as was the larger amount and the horizon. The principal was varied between $16 and $83, the later amount was varied between $30 and $85, and the horizon was varied between 10 and 70 days. This design yielded choices that had annualized discount rates between 128% and 1.2E+13. The average subject started switching over to the later option when discount rates were above 596%, and only 12% of subjects accepted the lowest discount rate of 128%. It is almost irrelevant if discounting over this horizon was constant or not, since one can undertake though experiments with longer horizons and rule out such rates being accepted for those choice.

One obvious concern with this task is that it is "almost hypothetical." Questionnaires were sent out to every undergraduate student at Williams College, which had a student population of roughly 2000 at the time. Subjects were told that the questionnaires could be returned by one of two days, and that on each day one would be drawn at random and one of the 21 choices played out for real. Despite cheap talk in the instructions that, to "make sure that you get a reward you prefer, you should *assume* that you are the winner, and then make *each choice* as though it were the one that you will win," these are very poor financial incentives. Even if the subject was certain that the lottery with the (delayed) $85 payment would be chosen, this is an expected earning of only $0.085. Although subjects were not told how many of the questionnaires were distributed, this is a small campus and such things are not private. In the event, 672 responded, implying that there was actually an expectation of only $0.25 if the largest prize was then selected. Moreover, even if the largest payment was chosen in all 21

---

[50] They also report (p.133) a small follow-up experiment with 5 subjects and a front end delay of two weeks. In that case 4 of the 5 choices entailed a switch from the later payment to the earlier payment.

[51] He also reports an experiment in which a willingness to accept a sooner payment in compensation for the subject's post-dated check for $50 in one month. This experiment suffered from even more serious credibility problems than the willingness to pay experiment (p.320) and will not be discussed.

cases, the expected earnings were only $56.19 per lucky subject, and not $85. It is a pity that this clean, transparent task was marred by the use of such poor incentives.

Kirby [1997] is a remarkable study: it used real incentives, used payments by subjects out of their own cash, used an incentive-compatible second-price sealed-offer auction to elicit present values, considered the effect of varying the deferred amount ($10 or $20), and considered all odd-numbered horizons between 1 and 29. Each subject entered 30 bids, and was told that one of these bids would be selected at random for payment if the bid was the winning bid. Each auction apparently consisted of the entire sample in an experiment, which does not affect the incentive compatibility of the procedure. Subjects in experiment 1 were "pseudo-volunteers" receiving extra credit in a psychology class for attending, but apart from the show-up rewards all payments were salient. Subjects in experiments 2 and 3 were "people from the Williams College community, including summer students, college staff, and persons unaffiliated with the college," and recruitment was by sign-up fliers and newspaper advertisements.

C. Recent Experimental Studies

Benhabib, Bisin and Schotter [2010] present subjects with two types of matching tasks. In one type, the subject was asked 30 questions of the form "what amount of money, $x, if paid to you today would make you indifferent to $y paid to you in t days?" In this case the amount $y and the horizon t would be filled in: y ∈ {10, 20, 30, 50, 100} and t ∈ {3 days, 1 week, 2 weeks, 1 month, 3 months, 6 months}. The response $x was incentivized with a Becker, DeGroot and Marschak [1964] (BDM) auction for one of the 30 choices selected at random. A price would be drawn from a uniform distribution between [$0, $y], and if the random price was greater than the stated amount $x then the subject would receive that random price immediately; otherwise the subject would receive $y in t days. So the upper bound of the BDM auction was the larger amount to be provided in the future. The other type of matching question involved 30 questions of the form, "what amount of money, $y, would make you indifferent between $x today and $y t days from now? [upper bound = $z]," where the text in brackets was given to subjects as notation instead of these words. In this case the values of t were the same as the first matching task, and the values of x ∈ {1, 2, 3, 5, 6, 7} for z=10, x ∈ {4, 7, 8, 10, 12, 14} for z=20, x ∈ {8, 14, 17, 19, 22, 24} for z=30, x ∈ {15, 20, 28, 32, 36, 39} for z=50, and x ∈ {40, 60, 65, 70, 75, 80} for z=100. The same subjects were given both sets of questions on different days. The data were evaluated using the flexible FC specification introduced in this study, and the model estimated for each individual using non-linear least squares. The individual estimates are very erratic, with a wide range of behaviors being inferred. The general theme is of extremely high discount rates, support for the fixed cost specification, and considerable noise.

Laury, McInnes and Swarthout [2012] build in controls for risk neutrality into the elicitation task, by asking subjects to make time-delay choices over binary lotteries. Since the binary lotteries are each defined over the same low prize and high prize, one can normalize the utility of each to 0 and 1 and then translate the probability of the lottery directly into a utility number (e.g., a 26.3% chance of the high prize has a utility value of 0.263, and a 55.5% chance of the high prize has a utility value of 0.555). This binary lottery method has been widely used in experimental economics for many years, and was first employed as such by Roth and Malouf [1979]. Their procedure offers subjects a multiple price list of choices between sooner and later options. The sooner option is always a 50:50 chance of getting $0 or $200 in 3 weeks time. The latter option offers increasing chances of the $200 in 12 weeks time. In the first row it is a 50% chance; in the second row it is 50.1%, in the third row it is 50.2%, and so on up to 64.9% in row 20.[52] Of course, these are strikingly similar choices from row to row, by any intuitive metric: in the middle of the table the expected value of the later option is varying from row to row by only 20 cents or 40 cents, compared to a sooner option with a constant expected value of $100.

Although behaviorally challenging for subjects, the elegance of this design is that subjects are choosing directly over time delays of utilities, so in a theoretical sense one can directly infer utility-adjusted discount rates. Of course, EUT is a maintained assumption of this approach. Their maximum-likelihood estimate of the discount rate is 12.2%, with a 95% confidence interval between 4.6% and 19.7% (their Table 7). They also undertake experiments with the same subjects to measure risk aversion and discounting over monetary flows, and infer a discount rate of 14.1%, with a 95% confidence interval between 6.6% and 21.5%, using the maximum-likelihood methods of AHLR [2008a].[53] These two maximum-likelihood estimates are not statistically different from each other.

Andreoni and Sprenger [2012a] also propose a "one shot" elicitation procedure that can control for non-linear utility and discounting at the same time, rather than requiring several procedures and joint estimation over the choices in those procedures to infer discount rates. Each subject is given a series of choices in which they allocate 100 fictitious tokens between a sooner option and a later option. The exchange rate between tokens and money is always fixed at $0.20 per token for the later option. For the sooner option the exchange rate varies from row to row.[54] In effect, this generates a convex budget set for individuals to choose from, avoiding the problems of indeterminacy with the linear budget set in the standard multiple price list identified by Cubitt and Read [2007]. It also allows one to estimate utility functions and discounting functions from the same set of choices, using parametric functional forms for both and assuming the validity of EUT.

Using non-linear least squares methods, Andreoni and Sprenger [2012a] estimate discount rates to be 29.8% in their preferred specification (their Table 1, specification 1). This overall estimate includes choices with no FED and choices with a FED of 7 days or 35 days: the estimated discount rate with no FED is 28.3%, and the estimated discount rates with the FED of 7 and 35 days are 32.9% and 26.7%, respectively (their Table 3, specification 1). All of these estimates have standard errors of around 6 percentage points, so all are statistically indistinguishable. Hence there is striking evidence here of non-hyperbolic discounting, albeit at very high rates. The estimates with no FED are virtually

---

[52] Since it matters for the interpretation of responses, we note that the increments from row to row are not constant. In percentage points, they are 0.1, 0.1, 0.2, 0.1, 0.2, 0.2, 0.2, 0.1, 0.2, 0.2, 0.2, 0.2, 0.2, 0.5, 0.6, 0.9, 2.4, 2.5 and 5.3, in rows 2 through 20, respectively. So subjects at row 14 and deciding whether to switch one row later or not will have a disproportionate impact on the average discount rate compared to subjects contemplating the same switch in earlier rows. Roughly 30% of their subjects were in this upper region (per their Figure 1; their Table 6 is a less reliable guide, since it excludes 26 of the 103 subjects due to their "re-switching" behavior).

[53] The monetary discounting task had information on the annual interest rate and the annual effective interest rate implied by each option, following Coller and Williams [1999]. The utility discounting task did not have this information (the decision sheet was on paper, and already had a lot of information).

[54] Their computer interface makes the implications of these choices clearer to subjects than they might seem from this description. The top of the display shows a calendar with the sooner and later dates displayed, and the bottom of the display instantly updates the actual amounts to be received sooner and later as the token allocation is changed by the subjects.

identical to those implied by the comparable treatment in Coller and Williams [1999], but the estimates with a FED are about a half of those of Coller and Williams [1999] from comparable treatments. The estimates with a FED are virtually identical to the estimates of Harrison, Lau and Williams [2002]. On the other hand, the estimates of Coller and Williams [1999] and Harrison, Lau and Williams [2002] make no effort to correct for non-linear utility, and those of Andreoni and Sprenger [2012a] do. The comparable, utility-adjusted estimates of AHLR [2008a] and Laury, McInnes and Swarthout [2012] *with a FED* are significantly lower, at 10.1% and 14.1% respectively. The comparable, utility-adjusted estimates of Coller, Harrison and Rutström [2012; Table 2, Panel A] *without a FED* are virtually identical, at 29.8%.

The econometric methods of Andreoni and Sprenger [2012a] have a significant flaw: the manner in which portfolio extremes are handled, when inferring the risk attitudes and discount rate of the representative agent.[55] The vast majority of their choices involve the subject choosing 0% or 100%, and indeed both extremes are observed. Elementary economic theory informs us that if the representative agent has a strictly concave utility function, even a modestly concave function, then *every* choice at *every* interest rate must be an interior choice. The only theoretical explanation for the vast majority of choices being at both extremes is that the representative agent has weakly convex utility functions. If one uses (linear or non-linear) ordinary least squares on these data, the model will seek to "fit" the average well, indicating a concave utility function since there are observed choices at both extremes, and the average is, well, in the interior. But this is a misleading inference, and in fact qualitatively wrong. The finding that subjects are risk neutral or risk lovers contrasts with *all* of the evidence for comparable samples, as reviewed by Harrison and Rutström [2008a]. It is not obvious if this is a reflection of the preferences of their particular college-student sample, an implication of subject confusion with their procedures, their misleading estimation methods, or all of these factors.

The inferences of Andreoni and Sprenger [2012b] suffer from a very different flaw. The basic design of Andreoni and Sprenger [2012a], which provides the control experiment for the claims in Andreoni and Sprenger [2012b], allows subjects in each choice to choose a portfolio of sooner payoffs and later payoffs. As noted above, they hold constant the later payoff amount, and vary the sooner payoff amount, but the idea is the same as our approach in the discount task, with the addition of allowing a portfolio allocation. For instance, using the "Sample Decision Sheet" in their Figure 1 (p.3340), and ignoring portfolio allocations for the moment, the first choice is between \$20 on a sooner date and \$20 on a later date, the second choice is between \$19 on the same sooner date and \$20 on the same later date, and so on to the last choice between \$14 on the sooner date and \$20 on the later date.[56] For each choice the subject decides on a fraction of their portfolio allocated to the sooner option and the residual fraction allocated to the later option. If the fraction for the sooner option in the first choice is 0.83, as in their example, the subject would receive \$16.60 = \$20 × 0.83 on the sooner date and \$3.40 = \$20 × 0.17 on the later date. If the fraction for the sooner option in the

[55] The online appendices of Andreoni and Sprenger [2012a] contain an elaborate series of efforts to show that different ways of econometrically handling the corner observations make no difference to their conclusions. None of these alternatives addresses the basic point about the economic theory needed to explain corner behavior, and in effect treats these corner observations as outside the structural model of risk preferences and discount rates. These arguments are presented in greater detail in Harrison, Lau and Rutström [2013].

[56] Their subject interface presents these amounts as "exchange rates" between 100 tokens and cash on those dates. So the second choice is actually presented as 100 tokens at an exchange rate of \$0.19 per token or 100 tokens at an exchange rate of \$0.20 per token. It is not clear why the added layer of tokens is needed.

second choice is 0.51, as in their example, the subject would receive $9.69 = $19 × 0.51 on the sooner date and $9.80 = $20 × 0.49 on the later date.

Andreoni and Sprenger [2012a] find that 70% of the observed choices in this design were at one extreme or the other, where the subject chose to allocate everything to the sooner payoff or the later payoff. They define the utility function over a single attribute, money, and find that the parameter that measures the concavity of the utility function is close to one, implying a utility function that is almost linear but still slightly concave.

The design of Andreoni and Sprenger [2012b] is to allow the sooner payoffs to be received with a probability p and the later payoffs to be received with probability P. They consider six {p, P} combinations: {1, 1}, {0.5, 0.5}, {0.5, 0.4}, {0.4, 0.5}, {1, 0.8} and {0.8, 1}. The initial, non-stochastic {1, 1} case is the design in Andreoni and Sprenger [2012a]. For the other treatments there was an *independent realization* of p and P. These realizations occurred at the same time, so that no preferences for temporal resolution of risk would confound inferences. The key finding is that more subjects allocated their portfolio in the interior when $p<1$ and/or $P<1$. Specifically, in their baseline {0.5, 0.5} treatment only 25% of the choices were at the extremes, compared to 74% for their {1, 1} control. The conclusion they draw from finding is that the *instantaneous* utility functions over risk and time are not the same.

The implication for the conclusion that "risk preferences are not time preferences" is immediate. If the intertemporal utility function that subjects in Andreoni and Sprenger [2012b] use is actually *non-additive*, then risk preferences over time streams of money need to be sharply distinguished from risk preferences over time streams over atemporal payoffs. In effect, there are two possible types of risk aversion when one considers risky choices over time, not one. The risk preferences over time streams has been called "correlation aversion" by Epstein and Tanny [1980], and is widely employed in the literature (e.g., Bommier [2007] and Denuit, Eeckhoudt and Rey [2010]). Risk preferences over atemporal payoffs can be called atemporal risk aversion. If one gives subjects choices over differently-time-dated payoffs, which is what Andreoni and Sprenger [2012b] did, *one sets up exactly the thought experiment that defines correlation aversion.* They compare behavior when subjects make choices over time-dated payoffs that are not stochastic with choices that subjects make over time-dated payoffs that are stochastic, and observe different choices. In the former case virtually all choices in their portfolios were at extreme allocations, either all payoffs sooner or all payoffs later; in the latter case they observed more choices in which subjects picked an interior mix of sooner and later payoffs, diversifying intertemporally. Evidence that subjects behave differently, when there is an opportunity for correlation aversion to affect their choices compared to a setting in which it has no role, is evidence of correlation aversion. It is not necessarily evidence for the claim that there is a "different instantaneous utility function" at work when considering stochastic and non-stochastic choices. We do not rule that hypothesis out, but there is a simpler explanation well within received theory. And, of course, there is considerable direct evidence for the empirical plausibility of correlation aversion in Andersen, Harrison, Lau and Rutström [2011].

Cheung [2014; Appendix A.2] proves nicely that correlation aversion provides an immediate explanation for the observed behavior in Andreoni and Sprenger [2010b].[57] He refers to it as a motive

[57] We first raised this point with Andreoni and Sprenger in a conference in Denmark in June 2010, and explained it to Cheung in 2011.

for intertemporal diversification, which is a valid way of characterizing the idea of correlation aversion. Hence, when Andreoni and Sprenger [2012b] claim that "risk preferences are not time preferences," one can restate this more carefully as "a-temporal risk aversion is not the same as intertemporal risk aversion," and of course that is correct whenever there is a non-additive intertemporal utility function.

Coller, Harrison and Rutström [2012] extend the procedures of Coller and Williams [1999] to focus on time delay choices over money that have no FED, and that have horizons varying between 1 day and 60 days. They also pool data on lottery choices from a sample of subjects drawn from the same population, and presented in Harrison, Johnson, McInnes and Rutström [2005]. With these data they estimate exponential and quasi-hyperbolic models using maximum likelihood, and then a mixture specification of the two.[58] They estimate the mixture weight on the exponential model at 0.59, and hence at 0.41 for the quasi-hyperbolic model; they cannot reject the hypothesis that this weight is 0.5, although they can reject the hypothesis that it is 0 or 1. In the mixture specification they estimate a value of 0.94 for the β of the quasi-hyperbolic model, and a value of δ for both models of 0.116. These estimates imply discount rates overall that range from thousands of percent for horizons of less than a week, around 1000% for a horizon of one week, around 200% for a horizon of two weeks, 63% for a horizon of a month, and 33% for a horizon of two months. Eventually these estimated rates would asymptote to 11.6%, but clearly there is evidence for hyperbolic discounting. Whether the weight of 0.41 on the quasi-hyperbolic model represents 41% of the sample or 41% of the choices of each subject, or something in between, is not clear from the estimates, although formally their mixture model assumed that it reflected each choice by each subject.

D. Estimation with Naturally Occurring Data

There have been several attempts to measure discounting functions using naturally occurring, non-experimental data. One that pays attention to the importance of jointly estimating the utility function and discount rates is Laibson, Repetto and Tobacman [2007]. They correctly note that experimental data and naturally occurring data should be viewed as complementary, and each has strengths and weaknesses. They use data from the United States, and a wide range of assumptions to estimate a quasi-hyperbolic and exponential discounting functions. In particular, in their baseline, preferred estimates they parametrically assume a CRRA of 2. This assumption leads them to find striking evidence in favor of quasi-hyperbolic behavior, with β estimated to be 0.70 and δ (in our notation) estimated to be 0.043, implying "short-run" discount rates for one-year horizon of 39.5% and long-run discount rates of 4.3%. Under these assumptions the exponential discount rate is 16.7%.

They report some sensitivity analyses of parameters (Table 4B), and note that when the CRRA is lowered to 1 the implied discount rates move closer to being exponential, although one cannot reject the hypothesis of hyperbolicky discounting. The short-run discount rate is then 24% and the long-run discount rate is 4%, with an exponential discount rate at 11.4%. Indeed, when they jointly estimate the RRA and discounting functions, as advocated by AHLR [2008a], the RRA point estimate is lowered to 0.22 for the quasi-hyperbolic model and 0.28 for the exponential model. These changes

[58] If one assumes that an exponential model characterizes the data, then their estimate of δ is 28.6% with a 95% confidence interval between 16.7% and 40.5% (their Table 2). If one assumes that a quasi-hyperbolic model characterizes the data, they estimate β to be 0.987 and δ to be 9.3%; a test that β=1 has a two-sided *p*-value of only 0.002.

in RRA understandably imply a short-run discount rate of 14.6%, with a long-run discount rate of 3.9% and an exponential discount rate of 9.2%. It is a pity that the bulk of their sensitivity analyses use the RRA of 2, given the clear importance of this preference parameter for the estimates. In one case they interact a parametrically assumed RRA equal to the estimated values, and other robustness checks (their "Compound Case D") and infer a short-run discount rate of 11% with a long-run discount rate of 5.6% and an exponential discount rate at 8.2%.

A crude evaluation of the uncertainty of these estimates can be inferred from the reported standard errors on the point estimates.[59] To take just the case of direct comparison to our results, where their preferred calibrating assumptions are used but the RRA is jointly estimated, the 90% confidence intervals for the short-run discount rate are 11.3% ↔ 18.1%, for the long-run discount they are 3.5% ↔ 4.3%, and for the exponential discount they are 7.6% ↔ 10.7%. This still generates evidence of hyperbolicky discounting, albeit in a less spectacular manner than their preferred specification with the RRA set parametrically to 2.

Additional References

Becker, Gordon M.; DeGroot, Morris H., and Marschak, Jacob., "Measuring Utility By A Single-Response Sequential Method," *Behavioral Science*, 9, July 1964, 226-232.

Bleichrodt, Han; Rohde, Kirsten I.M., and Wakker, Peter P., "Non-hyperbolic Time Inconsistency," *Games and Economic Behavior*, 66, 2009, 27-38.

Bommier, Antoine, "Risk Aversion, Intertemporal Elasticity of Substitution and Correlation Aversion," *Economics Bulletin*, 4(29), 2007, 1-8

Cheung, Stephen L., "On the Elicitation of Time Preferences Under Risk," *American Economic Review*, forthcoming 2014.

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

Denuit, Michel; Eeckhoudt, Louis, and Rey, Béatrice, "Some Consequences of Correlation Aversion in Decision Science," *Annals of Operations Research*, 176, 2010, 259-269.

Epstein, Larry G., and Tanny, Stephen M., "Increasing Generalized Correlation: a Definition and Some Economic Consequences," *Canadian Journal of Economics*, 13, 1980, 16-34.

Harrison, Glenn W.; Lau, Morten I., and Rutström, E. Elisabet,"Identifying Time Preferences with Experiments: Comment," *Working Paper*, Department of Risk Management and Insurance,

[59] We generate 100,000 random normal deviates with the mean set equal to the reported point estimates and the standard deviation set equal to the reported standard error. We then generate the discount rates using the non-linear formulae appropriate for their specification (Laibson, Repetto and Tobacman [2007; p.11]), and report the confidence intervals of these generated discount rates. The correct calculation would use the "delta method" to correctly infer the standard errors for the discount rates (Oehlert [1992]), but for that one needs access to the estimated covariance matrix.

Robinson College of Business, Georgia State University, February 2013.

Harvey, Charles M., "Value Functions for Infinite-Period Planning," *Management Science*, 32(9), 1986, 1123-1139.

Harvey, Charles M., "The Reasonableness of Non-Constant Discounting," *Journal of Public Economics*, 53, 1994, 31-51.

Herrnstein, R.J., "Relative and Absolute Strength of Response as a Function of Frequency of Reinforcement," *Journal of the Experimental Analysis of Behavior*, 4, 1961, 267–272.

Jamison, Dean T., and Jamison, Julian, "Characterizing the Amount and Speed of Discounting Procedures," *Journal of Benefit-Cost Analysis*, 2(2), April 2011, Article 1.

Killeen, Peter R., "An Additive-Utility Model of Delay Discounting," *Psychological Review*, 116, 2009, 602-619.

Kim, B. Kyu, and Zauberman, Gal, "Perception of Anticipatory Time in Temporal Discounting," *Journal of Neuroscience, Psychology and Economics*, 2(2), 2009, 91-101.

Laibson, David; Repetto, Andrea, and Tobacman, Jeremy, "Estimating Discounting functions with Consumption Choices over the Lifecycle," *Working Paper*, Department of Economics, Harvard University, November 2007.

Mazur, James E., "An Adjustment Procedure for Studying Delayed Reinforcement," in M.L. Commons, J.E. Mazur, J.A. Nevin and H. Rachlin (eds.), *The Effect of Delay and Intervening Events on Reinforcement Value* (Hillsdale, NJ: Erlbaum, 1987).

Oehlert, Gary W., "A Note on the Delta Method," *American Statistician*, 46(1), February 1992, 27-29.

Roelofsma, Peter H.M.P., "Modelling Intertemporal Choices," *Acta Psychologica*, 93, 1996, 5-22.

Roth, Alvin E., and Malouf, Michael W. K., "Game-Theoretic Models and the Role of Information in Bargaining," *Psychological Review*, 86, 1979, 574-594.

Scholten, Marc, and Read, Daniel, "Discounting By Intervals: A Generalized Model of Intertemporal Choice," *Management Science*, 52, 2006, 1426-1438.

Takahashi, T., "Loss of Self-Control in Intertemporal Choice May Be Attributable to Logarithmic Time Perception," *Medical Hypotheses*, 65(4), 2005, 691-693.

Zauberman, Gal; Kim, B. Kyu; Malkoc, Selin A., and Bettman, James R., "Time Discounting and Discounting Time: Subjective Time Preference and Intertemporal Preferences," *Journal of Marketing Research*, 46, 2009, 543-556.

Appendix E: Econometric Specification (WORKING PAPER)

A. Estimating the Utility Function

Assume for the moment that utility of income is defined by

$$U(y) = M^{(1-r)}/(1-r) \tag{1}$$

where M is the lottery prize and r≠1 is a parameter to be estimated. For r=1 assume U(M)=ln(M) if needed. Thus r is the coefficient of CRRA: r=0 corresponds to risk neutrality, r<0 to risk loving, and r>0 to risk aversion. Let there be two possible outcomes in a lottery. Under EUT the probabilities for each outcome M_j, $p(M_j)$, are those that are induced by the experimenter, so expected utility is simply the probability weighted utility of each outcome in each lottery i plus some level of background consumption ω:

$$EU_i = [\ p(M_1) \times U(\omega+M_1)\ ] + [\ p(M_2) \times U(\omega+M_2)\ ] \tag{2}$$

The EU for each lottery pair is calculated for a candidate estimate of r, and the index

$$\nabla EU = EU_R - EU_L \tag{3}$$

calculated, where EU_L is the "left" lottery and EU_R is the "right" lottery as presented to subjects. This latent index, based on latent preferences, is then linked to observed choices using the cumulative logistic distribution function Λ(∇EU). This "logit" function takes any argument between ±∞ and transforms it into a number between 0 and 1. Thus we have the logit link function,

$$\text{prob(choose lottery R)} = \Lambda(\nabla EU) \tag{4}$$

The index defined by (3) is linked to the observed choices by specifying that the R lottery is chosen when Λ(∇EU)>½, which is implied by (4).

Thus the likelihood of the observed responses, conditional on the EUT and CRRA specifications being true, depends on the estimates of r given the above statistical specification and the observed choices. The conditional log-likelihood is then

$$\ln L(r; y, \omega, \mathbf{X}) = \textstyle\sum_i [\ (\ln \Lambda(\nabla EU)\times \mathbf{I}(y_i = 1)) + (\ln (1\text{-}\Lambda(\nabla EU))\times \mathbf{I}(y_i = -1))\ ] \tag{5}$$

where $\mathbf{I}(\cdot)$ is the indicator function, $y_i = 1(-1)$ denotes the choice of the Option R (L) lottery in risk aversion task i, and **X** is a vector of individual characteristics reflecting age, sex, race, and so on. The parameter r is defined as a linear function of the characteristics in vector **X**.

Harrison and Rutström [2008a; Appendix F] review procedures and syntax from the popular statistical package *Stata* that can be used to estimate structural models of this kind, as well as more complex non-EUT models. The goal is to illustrate how experimental economists can write explicit maximum likelihood (ML) routines that are specific to different structural choice models. It is a simple matter to correct for stratified survey responses, multiple responses from the same subject ("clustering"), or heteroskedasticity, as needed.

Extensions of the basic model are easy to implement, and this is the major attraction of the structural estimation approach. For example, one can easily extend the functional forms of utility to allow for varying degrees of relative risk aversion (RRA). Consider, as one important example, the Expo-Power (EP) utility function proposed by Saha [1993]. Following Holt and Laury [2002], the EP function is defined as

$$U(x) = [1-\exp(-\alpha x^{1-ŕ})]/\alpha, \tag{1'}$$

where α and ŕ are parameters to be estimated. RRA is then $ŕ + \alpha(1-ŕ)y^{1-ŕ}$, so RRA varies with income if $\alpha \neq 0$. This function nests CRRA (as $\alpha \rightarrow 0$) and CARA (as $ŕ \rightarrow 0$).

It is also simple matter to generalize this ML analysis to allow the core parameter r to be a linear function of observable characteristics of the individual or task. We would then extend the model to be $r = r_0 + R \times \mathbf{X}$, where r_0 is a fixed parameter and R is a vector of effects associated with each characteristic in the variable vector **X**. In effect the unconditional model assumes $r = r_0$ and just estimates r_0. This extension significantly enhances the attraction of structural ML estimation, particularly for responses pooled over different subjects, since one can condition estimates on observable characteristics of the task or subject.

An important extension of the core model is to allow for subjects to make some errors. The notion of error is one that has already been encountered in the form of the statistical assumption that the probability of choosing a lottery is not 1 when the EU of that lottery exceeds the EU of the other lottery. This assumption is clear in the use of a link function between the latent index ∇EU and the probability of picking one or other lottery; in the case of the logistic CDF, this link function is $\Lambda(\nabla EU)$. If there were no errors from the perspective of EUT, this function would be a step function: zero for all values of $\nabla EU<0$, anywhere between 0 and 1 for $\nabla EU=0$, and 1 for all values of $\nabla EU>0$.

The problem with this CDF is immediate: it predicts with probability one or zero. The likelihood approach asks the model to state the probability of observing the actual choice, conditional on some trial values of the parameters of the theory. Maximum likelihood then locates those parameters that generate the highest probability of observing the data. For binary choice tasks, and independent observations, the likelihood of the sample is just the product of the likelihood of each choice conditional on the model and the parameters assumed, and that the likelihood of each choice is just the probability of that choice. So if we have any choice that has zero probability, and it might be literally 1-in-a-million choices, the likelihood for that observation is not defined. Even if we set the probability of the choice to some arbitrarily small, positive value, the log-likelihood zooms off to minus infinity. We can reject the theory without even firing up any statistical package.

Of course, this implication is true for any theory that predicts deterministically, including Expected Utility Theory. This is why one needs some formal statement about how the deterministic prediction of the theory translates into a probability of observing one choice or the other, and then perhaps also some formal statement about the role that structural errors might play.[60] In short, one

[60] Exactly the same insight in a strategic context leads one from Nash Equilibria to Quantal Response Equilibria, if one re-interprets the CDF in terms of best-response functions defined over expected (utility) payoffs from two strategies. The only difference in the maximum likelihood specification is that the equilibrium condition jointly constrains the likelihood of observing certain choices by two or more players.

cannot divorce the job of the theorist from the job of the econometrician, and some assumption about the process linking latent preferences and observed choices is needed. That assumption might be about the mathematical form of the link, as in (3), but it cannot be avoided. Even the very definition of risk aversion needs to be specified using stochastic terms unless we are to impose absurd economic properties on estimates (Wilcox [2008][2011]).

We employ the error specification originally due to Fechner and popularized by Hey and Orme [1994]. This error specification posits the latent index

$$\nabla EU = (EU_R - EU_L)/\mu \tag{3$'$}$$

instead of (3), where μ is a structural "noise parameter" used to allow some errors from the perspective of the deterministic EUT model. This is just one of several different types of error story that could be used, and Wilcox [2008] provides a masterful review of the implications of the alternatives.[61] As $\mu \to 0$ this specification collapses to the deterministic choice EUT model, where the choice is strictly determined by the EU of the two lotteries; but as μ gets larger and larger the choice essentially becomes random. When $\mu=1$ this specification collapses to (3), where the probability of picking one lottery is given by the ratio of the EU of one lottery to the sum of the EU of both lotteries. Thus μ can be viewed as a parameter that flattens out the link functions as it gets larger.

An important contribution to the characterization of behavioral errors is the "contextual error" specification proposed by Wilcox [2011]. It is designed to allow robust inferences about the primitive "more stochastically risk averse than." It posits the latent index

$$\nabla EU = ((EU_R - EU_L)/\nu)/\mu \tag{3$''$}$$

instead of (3$'$), where ν is a new, normalizing term for each lottery pair L and R. The normalizing term ν is defined as the maximum utility over all prizes in this lottery pair minus the minimum utility over all prizes in this lottery pair. The value of ν varies, in principle, from lottery choice to lottery choice: hence it is said to be "contextual." For the Fechner specification, dividing by ν ensures that the *normalized* EU difference $[(EU_R - EU_L)/\nu]$ remains in the unit interval.

The likelihood of the risk aversion task responses, conditional on the EUT and CRRA specifications being true, depends on the estimates of r and μ. The conditional log-likelihood is

$$\ln L(r, \mu; y, \mathbf{X}) = \sum_i [(\ln \Lambda(\nabla EU)\times \mathbf{I}(y_i=1)) + (\ln (1-\Lambda(\nabla EU))\times \mathbf{I}(y_i=-1))] \tag{6}$$

where $y_i = 1(-1)$ denotes the choice of Option B (A) in risk aversion task i, and $\mathbf{X}$ is a vector of individual characteristics. The value of ν depends on the data, and not on the estimated preference parameters r and μ.

---

[61] Some specifications place the error at the final choice between one lottery or after the subject has decided which one has the higher expected utility; some place the error earlier, on the comparison of preferences leading to the choice; and some place the error even earlier, on the determination of the expected utility of each lottery.

B. Estimating the Discounting Function

Assume EUT holds for choices over risky alternatives and that discounting is exponential. A subject is indifferent between two income options M_t and $M_{t+\tau}$ if and only if

$$(1/(1+\delta)^t)\ U(\omega+M_t) + (1/(1+\delta)^{t+\tau})\ U(\omega) = (1/(1+\delta)^t)\ U(\omega) + (1/(1+\delta)^{\tau})\ U(\omega+M_{t+\tau}) \qquad (7)$$

where $U(\omega+M_t)$ is the utility of monetary outcome M_t for delivery at time t plus some measure of background consumption ω, δ is the discount rate, τ is the horizon for delivery of the later monetary outcome at time $t+\tau$, and the utility function U is separable and stationary over time. The left hand side of equation (6) is the sum of the discounted utilities of receiving the monetary outcome M_t at time t (in addition to background consumption) and receiving nothing extra at time $t+\tau$, and the right hand side is the sum of the discounted utilities of receiving nothing over background consumption at time t and the outcome $M_{t+\tau}$ (plus background consumption) at time $t+\tau$. Thus (6) is an indifference condition and δ is the discount rate that equalizes the present value of the *utility* of the two monetary outcomes M_t and $M_{t+\tau}$, after integration with an appropriate level of background consumption ω.

We can write out the likelihood function for the choices that our subjects made and jointly estimate the risk parameter r in equation (1) and the discount rate parameter δ in (6). We use the same stochastic error specification as in (3′), albeit with a different Fechner error term υ for the discount choices.[62] Instead of (3′) we have

$$\nabla PV = (PV_A - PV_B)/\eta, \qquad (8)$$

where the discounted utility of Option A is given by

$$PV_A = (1/(1+\delta)^t)(\omega+M_A)^{(1-r)}/(1\text{-}r) + (1/(1+\delta)^{t+\tau})\ \omega^{(1-r)}/(1\text{-}r) \qquad (9)$$

and the discounted utility of Option B is

$$PV_B = (1/(1+\delta)^t)\ \omega^{(1-r)}/(1\text{-}r) + (1/(1+\delta)^{t+\tau})\ (\omega+M_B)^{(1-r)}/(1\text{-}r), \qquad (10)$$

and M_A and M_B are the monetary amounts in the choice tasks presented to subjects. The parameter η captures noise for the discount rate choices, just as μ was a noise parameter for the risk aversion choices.[63] As noted in the text, we assume that the utility function is stable over time and is perceived *ex ante* to be stable over time.[64] We also assume that the parameter $r < 1$, to ensure that $\delta > 0$.

---

[62] We do not need to apply the contextual utility correction ν for these choices since they are over deterministic monetary amounts.

[63] It is not obvious that $\mu=\eta$, since these are cognitively different tasks. Our own priors are that the risk aversion tasks are harder, since they involve four outcomes compared to two outcomes in the discount rate tasks, so we would expect $\mu>\eta$. Error structures are things one should always be agnostic about since they capture one's modeling ignorance, and we allow the error terms to differ between the risk and discount rate tasks.

[64] Direct evidence for the former proposition is provided by Andersen, Harrison, Lau and Rutström [2008b], who examine the temporal stability of risk attitudes in the Danish population. The second proposition is a more delicate matter: even if utility functions are stable over time, they may not be subjectively perceived to be, and that is what matters for use to assume that the same r that appears in (1) appears in (9) and (10). When there

Thus the likelihood of the discount rate responses, conditional on the EUT, CRRA and exponential discounting specifications being true, depends on the estimates of r, δ, μ and η, given the assumed value of ω and the observed choices. The conditional log-likelihood is

$$\ln L\,(r, \delta, \mu, \eta; y, \omega, \mathbf{X}) = \textstyle\sum_i \left[(\ln \Lambda(\nabla PV) \times \mathbf{I}(y_i=1)) + (\ln (1-\Lambda(\nabla PV)) \times \mathbf{I}(y_i=-1)) \right] \tag{11}$$

where $y_i = 1(-1)$ again denotes the choice of Option B (A) in discount rate task i, and **X** is a vector of individual characteristics.

The joint likelihood of the risk aversion and discount rate responses can then be written as

$$\ln L\,(r, \delta, \mu, \eta; y, \omega, \mathbf{X}) = \ln L^{RA} + \ln L^{DR} \tag{12}$$

where L^{RA} is defined by (6) and L^{DR} is defined by (11). This expression can then be maximized using standard numerical methods. As explained in the main text, the parameter ω is set exogenously: using data from the household expenditure survey at Statistics Denmark, AHLR [2008a; p.600, Appendix D] calculate per capita consumption of private nondurable goods on an average daily basis as being equal to 118 kroner in 2003.[65] We adjust that amount for inflation to the time of our experiments, and assume ω = 130 kroner.

Nothing in this inferential procedure relied on the use of EUT, or the CRRA functional form. Nor did anything rely on the use of the E discounting function. These methods generalize immediately to alternative models of decision making under risk, and especially to alternative discounting functions.

Additional References

Saha, Atanu, "Expo-Power Utility: A Flexible Form for Absolute and Relative Risk Aversion," *American Journal of Agricultural Economics*, 75(4), November 1993, 905-913.

Wilcox, Nathaniel T., "Stochastic Models for Binary Discrete Choice Under Risk: A Critical Primer and Econometric Comparison," in J. Cox and G.W. Harrison (eds.), *Risk Aversion in Experiments* (Bingley, UK: Emerald, Research in Experimental Economics, Volume 12, 2008).

Wilcox, Nathaniel T., "'Stochastically More Risk Averse:' A Contextual Theory of Stochastic Discrete Choice Under Risk," *Journal of Econometrics*, 162(1), May 2011, 89-104.

---

is no front end delay, this assumption is immediate for (9), but not otherwise. But whether or not individuals suffer from a "projection bias" is a deep matter, demanding more research: see Ainslie [1992; p. 144-179, §6.3], Kirby and Guastello [2001] and Loewenstein, O'Donoghue and Rabin [2003].

[65] AHLR [2008a; p.602] show that estimates are robust to variations of ω between 50 and 200 kroner.

Appendix F: Robustness Checks (WORKING PAPER)

A. Probability Weighting

One popular alternative to EUT is to allow the decision-maker to transform the objective probabilities presented in lotteries and to use these weighted probabilities as decision weights when evaluating lotteries. If w(p) is the probability weighting function assumed, and one only has lotteries with two prizes, as here, then

$$EU_i = [\ p(M_1) \times U(\omega+M_1)\ ] + [\ p(M_2) \times U(\omega+M_2)\ ] \tag{1}$$

becomes

$$RDEU_i = [\ w(p(M_1)) \times U(\omega+M_1)\ ] + [\ (1\text{-}w(p(M_1))) \times U(\omega+M_2)\ ], \tag{1$'$}$$

where RDEU refers to the Rank-Dependent Expected Utility model of Quiggin [1982], and the remaining econometric specification remains the same. Of course, one then has to specify the functional form for w(p) and estimate additional parameters, but the logic extends naturally.

Prelec [1998] offers a two-parameter probability weighting function that exhibits considerable flexibility. This function is

$$w(p) = \exp\{-\eta(-\ln p)^{\varphi}\}, \tag{2}$$

and is defined for $0<p<1$, $\eta>0$ and $\varphi>0$. Indeed, when $\varphi=1$ this function collapses to the venerable power function

$$w(p) = p^{\eta} \tag{3}$$

The implied probability weighting function is shown in Figure F1, and it is S-shaped and crosses the diagonal at the midpoint. We thus see some evidence in favor of probability weighting, in which small (high) probabilities are weighted less (more) than under EUT. The implied discount rate, with the Exponential discounting specification, is 8.1% with a 95% confidence interval between 6.2% and 9.9%. Although this is not a statistically significant decrease compared to the comparable EUT specification (Table 1, Panel A), it directly reflects the fact that with probability weighting the utility function is further from being risk neutral: the estimate of r with probability weighting is 0.75 instead of 0.65 under EUT. As noted earlier, a direct application of Jensen's Inequality shows that a more concave utility function must decrease inferred discount rates for given choices between the two monetary options.[66]

[66] The dimension of risk preferences that is captured through the probability weighting function does not affect the choices across the discount rate tasks since no probabilities are present. The dimension of risk preferences that is modeled through the utility function implies more concavity of the utility function, so a slightly lower discount rate is required in order to explain the observed discounting choices.

B. Mixture Models

Mixture specifications allow two or more data-generating processes to explain observed behavior. Applications of this idea show clear evidence that behavior is not wholly explained by any one of the popular models of discounting behavior or decision making under risk. AHLR [2008; §3.D] consider a mixture specification of exponential and hyperbolic discounting, and find that 72% of the choices are better characterized as exponential. This estimate of the mixing probability is statistically significantly different from 0 or 50%. Similarly, Harrison and Rutström [2009] find roughly equal support for EUT and Prospect Theory in a lab setting; Harrison, Humphrey and Verschoor [2010] find roughly equal support for EUT and Rank-Dependent Utility models in artefactual field experiments in India, Ethiopia and Uganda; and Coller, Harrison and Rutström [2012] find roughly equal support for exponential and quasi-hyperbolic discounting in the laboratory.

The key insight from mixture specifications is to simply change the question that is posed to the data. Previous econometric analyses have posed a proper question: if one and only one data-generating process is to account for these data, what are the estimated parameter values and do they support a non-standard specification? The simplest, finite mixture specification changes this to: if two data-generating processes are allowed to account for the data, what fraction is attributable to each, and what are the estimated parameter values? So stated, one can imagine someone still wanting to ask the former question, if they just wanted one "best" model. But that question is also seen to constrain evidence of heterogeneity of decision-making processes, and we prefer to avoid that when we can. There are fascinating issues with the specific implementation and interpretation of mixture models, but those are not germane to the main insight they provide.[67]

The results are presented in Table F1 for two sets of mixture models. The first, in panel A, is a mixture of the Exponential and Quasi-Hyperbolic discounting model. The mixing probability, π^{E}, is estimated to be 0.276, so the complementary probability, π^{QH}, is 0.724. Thus it would appear that there is considerably more support for the QH specification, until one examines the estimated parameter values for each model. In the case of the Exponential model the discount rate is 15.3%, and for the Quasi-Hyperbolic model it is effectively a second Exponential specification with a discount rate of 3.3% because the estimate of β is essentially 1. Thus we see a bimodal distribution in the sample, with just over a quarter of the choices being characterized by a discount rate of 15.3% and the rest by a discount rate of 3.3%.

This pattern also arises when we consider a mixture of Exponential and Weibull discounting models, in panel B of Table F1. In this case the sample divides into one mode with 29% of choices at a discount rate of 14.7%, and the other mode with Weibull discounting that is not statistically significantly different from Exponential (*p*-value = 0.21). The discount rates implied by the Weibull parameter estimates range from 7.6% for a 1 day horizon, to 5.2% for a one week horizon, to 3.9%

[67] For example, does one constrain individuals or task types to be associated with just one data-generating process, or allow each choice to come from either? Does one consider more than two types of processes, using some specification rule to decide if there are 2, or 3, or more? Does one specify general models for each data-generating process and see if one of them collapses to a special case, or just specify the competing alternatives explicitly from the outset? How does one check for global maximum likelihood estimates in an environment that might generate multi-modal likelihood functions "naturally"? Harrison and Rutström [2009] discuss these issues, and point to the older literature.

for a one month horizon, down to 2.5% for a one year horizon. But it is important to recognize the relative statistical imprecision of the implied discount rates in this Weibull specification for the shorter horizons. Figure F2 illustrates this, showing the point estimates for horizons up to 3 months along with the 95% confidence intervals. It is easy to see why one cannot formally reject the assumption of constant discounting with this Weibull specification, and yet there is "movement" in discount rates as the horizon gets larger. Of course, the quantitative size of these discount rates pale in comparison to the hundreds of percent reported in some literature, but there is nonetheless the suggestion of a non-constant discounting agent struggling to be identified.

Table F2 allows us to see more clearly what factors affect the use of the Exponential and Weibull discounting models.[68] The mixing probability is allowed to vary with the treatment dummies, and we see significant effects from three of the treatments. The use of a front end delay increases the probability of the Exponential model being used by 14.7 percentage points. Thus the weight on the Exponential δ of 13.8% increases from 0.353 to 0.5 (= 0.353 + 0.147), and shifts to the Weibull specification. The provision of information on implied interest rates reduces the weight on the Exponential model by 14.0 percentage points, and the use of an increasing horizon in the presentation of the discounting task reduces weight on the Exponential model by 12.4 percentage points. The implied Weibull discount rates are virtually identical to those displayed in Figure F2. But the Weibull parameter s is estimated more precisely, and one can reject the assumption that the Weibull discounting function collapses to an Exponential discounting function with a *p*-value of 0.056. Again, these non-constant discount rates are modest in size compared to the conventional wisdom, but now they are approaching statistical significance.

In fact, a mixture model of two Exponential distributions provides a parsimonious characterization of the data. This model places a mixing probability of 0.28 on a discount rate of 15.1% and a mixing probability of 0.72 on a discount rate of 3.3%. The log-likelihood for this specification is -18511.1, only slightly worse than the two specifications in Table F1.

C. Observed Individual Heterogeneity

It is possible to condition our core parameters on individual demographic covariates, just as we considered covariates for treatment earlier. Table F3 contains the maximum likelihood estimates with the Exponential model. Unless otherwise noted, all variables are binary. Variable FEMALE indicates a female; YOUNG is someone aged less than 30; MIDDLE is someone aged between 40 and 50; OLD is someone aged over 50 (so the omitted age category are those aged between 30 and 39); EDUCATION is someone who has substantial higher education[69]; and HIGH INCOME is someone with household income in 2009 of 500,000 kroner or more.

The only demographic covariate to have any statistically significant impact on elicited discount rates is whether the individual is a female. Women have discount rates that are 4.3 percentage points lower than men, and the *p*-value on this estimated effect of 0.088. In turn, this derives from women being more risk averse: their RRA is 0.278 higher than men, with a *p*-value on this estimated effect of 0.018. Hence they have a more concave utility function and a lower implied discount rate. Looking at

[68] Comparable results are obtained if we extend the mixture of Exponential and Quasi-Hyperbolic, but the log-likelihood of that specification is inferior to the log-likelihood of the Exponential and Weibull mixture.

[69] Specifically, the completion of medium-cycle or longer-cycle higher education.

total effects instead of marginal effects, men on average have discount rates of 13.8% and women have discount rates of 10.0%, and the difference is statistically significant (*p*-value = 0.053).

D. Unobserved Individual Heterogeneity

We account for unobserved individual heterogeneity through the possibility that errors are clustered by the subject that the choices are associated with, but one can also allow unobserved individual heterogeneity in the population to be characterized by *random coefficients* following some parametric distribution. In other words, in the Exponential discounting model, one can allow the coefficients r and δ to be distributed in a random manner across the population: each subject behaves as if they have a specific r and δ, but there is variation across subjects and that variation is assumed to be characterized by some parametric distribution. If δ is assumed to vary according to a Normal distribution, then one would estimate two "hyper-parameters" to characterize that distribution: a population mean of δ and a population standard *deviation* of δ.[70] Each of these hyper-parameters would have a point estimate and a standard *error*, where the latter derives from familiar sampling variability. As the sample size increases, and assuming consistent estimators, the sampling error on the point estimate of the population mean and the point estimate of the population standard deviation would converge to 0, but there is no presumption that the point estimate for the population standard deviation converge to 0, since it is a characteristic of the population and not sample variability.

One extension which is required here is to allow for the latent index to be a non-linear function of core parameters, but we use methods developed by Andersen, Harrison, Hole, Lau and Rutström [2012], and reviewed in Appendix C (available online), to estimate such specifications. In fact, we also allow the distribution for r and δ to be a Logit-Normal distribution, which is a logistic transform of a normally distributed variable. Due originally to Johnson [1949], and familiar in bio-statistics, this transformation allows the resulting distribution to closely approximate a flexible Beta distribution: it allows skewness and bimodality. The domain is restricted to the unit interval, but it is a simple matter to expand that to any finite interval.

Figure F3 illustrates these estimates, for the base case specification (CRRA and Exponential discounting). In each case the parameters of the underlying Normal distribution are shown, and the logistic transform Λ then applied to them. For the risk aversion parameter r we estimate a mean for the Normal distribution of -0.23 and a standard deviation of 0.79: these are *not* the estimates for the parameter r itself, but the parameters defining the argument of the logistic transform, so we end up with the population distribution Λ(N(-0.23, 0.79)) shown in Figure F3. The population distribution is generally risk averse, with a mean of 0.56, a median of 0.55, a standard deviation of 0.18, and a skewness of -0.17 (where a symmetric distribution has a skewness of 0). The population distribution for the discount rate is sharply, positively skewed.[71] The average discount rate in the population is

---

[70] In fact we allow for a non-zero correlation between these two random coefficients, so their covariance is a third hyper-parameter to be estimated.

[71] There is one technical issue of importance here, however. As flexible as the Logit-Normal is, it only allows bimodality at the *end-points* of the finite interval allowed. In this case we constrained the domain to be between 0 and 0.6, and hence the mode close to 0 *might* be an artefacts of that assumption. Although we know a priori that $\delta \geq 0$, we do not know the upper bound. One can loop through alternative parametric assumptions of the upper bound and evaluate the maximum likelihood at each: these are known as profile likelihoods. In our case the qualitative results are invariant to assuming upper bounds lower than 0.6. A better solution, and common in

0.14, the median is 0.078, the standard deviation is 0.15, and the skewness is 1.19. Since this is a skewed distribution, one should not infer statistical insignificance from the standard deviation exceeding the mean. For the same reason, the appropriate measure of central tendency of the population distribution is the median rather than the mean. The covariance between the two random coefficients is -0.21, with a 95% confidence interval between -0.26 and -0.16; so we reject the hypothesis that the two coefficients are independent. This covariance implies a correlation of -0.13, which of course is consistent with the application of Jensen's Inequality, which shows that a *more* concave utility function must *decrease* inferred discount rates for given choices between the two monetary options.

Figure F4 displays comparable estimates of the two Quasi-Hyperbolic parameters. Consistent with the earlier estimates, there is no evidence for the instantaneous discounting premium that occurs when $\beta<1$. The population distribution for β is estimated to lie tightly around 1: the mean of the distribution is 0.999 and the median is 1.000. The estimated population distribution for δ in Figure F4 is similar to the estimates for the Exponential model in Figure F3, which is not surprising given that β is so close to 1.

E. Linear Utility

Table F4 replicates the estimates of Table 1 but with the assumption that the utility functions are linear. We find no evidence that this assumption changes our conclusion about the lack of evidence for hyperbolic discounting. The Exponential discounting model implies an 18.3% discount rate, with a 95% confidence interval between 15.5% and 21.2%, so this is lower than the estimate reported in AHLR [2008a] (25.2%, with a 95% confidence interval between 22.8% and 27.6%). The β-coefficient in the Quasi-Hyperbolic discounting model has a value of 1.005 and is not significantly different from 1 (*p*-value of 0.64), and there is no significant evidence of non-constant discounting in the Fixed-Cost and Weibull discounting models.

The Weibull estimates in Table F4 imply discount rates of 19.3% for a 1 week horizon, with a 95% confidence interval between 4.7% and 33.9%. After 2 weeks the estimated rate is 18.6% (7.7% ↔ 29.5%), after 1 month it is 17.9% (10.8% ↔ 25.0%), after 3 months it is 16.9% (14.2% ↔ 19.7%), and after 1 year it is 15.8% (11.1% ↔ 20.4%). These differences in predicted discount rates for different time horizons are not statistically significant.

the statistical literature to allow "internal modes," is to allow mixtures.

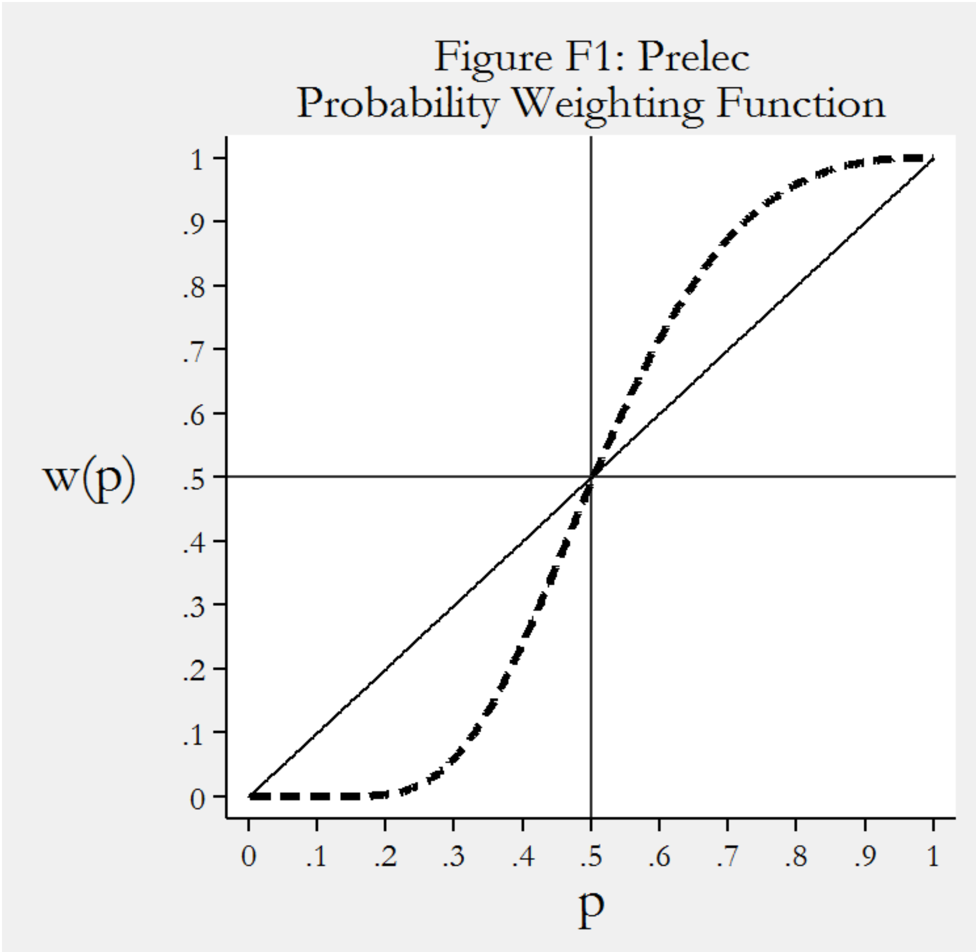
Figure F1: Prelec
Probability Weighting Function
w(p)
p
0
.1
.2
.3
.4
.5
.6
.7
.8
.9
1

Table F1: Estimates of Mixture Models

Parameter	Point Estimate	Standard Error	*p*-value	95% Confidence Interval	
A. Exponential and Quasi-Hyperbolic Discounting (LL = -18510.7)					
Exponential Model					
π^E	0.276	0.073	<0.001	0.133	0.419
δ	0.153	0.031	<0.001	0.092	0.213
Quasi-Hyperbolic Model					
π^{QH}	0.724	0.073	<0.001	0.581	0.87
β	1.001	0.003	<0.001	0.995	1.001
δ	0.033	0.007	<0.001	0.019	0.048
H_0: $\beta = 1$, *p*-value = 0.83					
B. Exponential and Weibull Discounting (LL = -18504.3)					
Exponential Model					
π^E	0.195	0.029	<0.001	0.138	0.253
δ	0.206	0.019	<0.001	0.109	0.184
Weibull Model					
π^W	0.71	0.039	<0.001	0.633	0.786
$\acute{r}$	0.03	0.004	<0.001	0.022	0.039
$\acute{s}$	1.233	0.188	<0.001	0.864	1.602
H_0: $\acute{s} = 1$, *p*-value = 0.21					

Figure F2: Weibull Discount Rates from Mixture Model

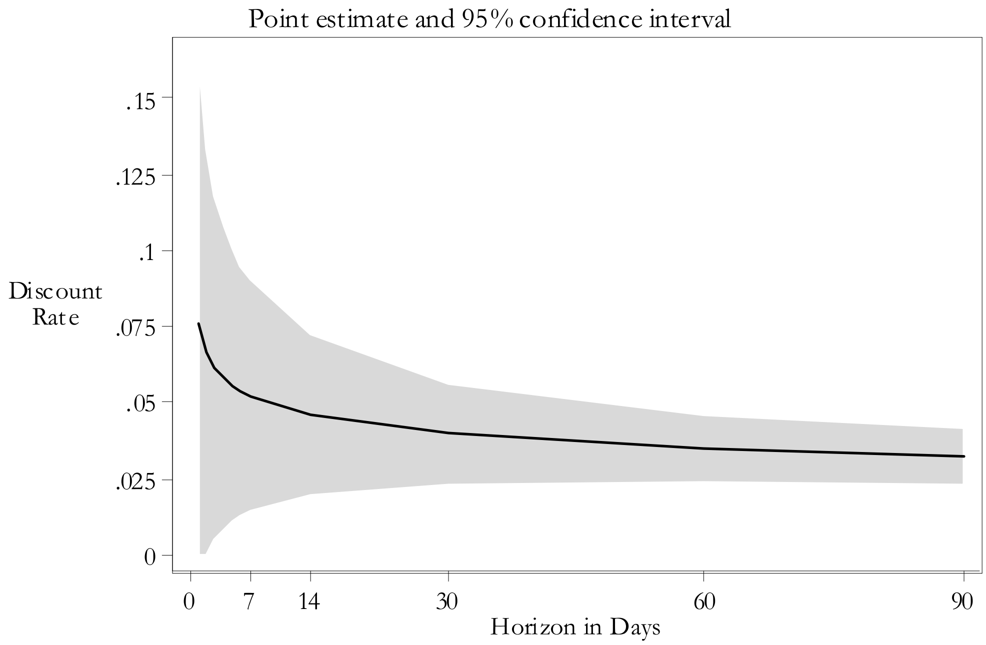

Table F2: Estimates of Exponential and Weibull Mixture Model
LL = -18358.9

Parameter	Point Estimate	Standard Error	*p*-value	95% Confidence Interval	
		Mixing Probability for Exponential Model			
π^E Constant	0.353	0.074	<0.001	0.209	0.498
FED	0.147	0.064	0.023	0.021	0.273
INFO	-0.14	0.055	0.012	-0.248	-0.031
H_ORDER	-0.124	0.059	0.035	-0.239	-0.009
P_HIGH	0.034	0.063	0.587	-0.089	0.159
RA_FIRST	0.052	0.067	0.435	-0.079	0.184
FEE_HIGH	0.002	0.062	0.978	-0.119	0.122
		Exponential Model			
δ	0.138	0.017	<0.001	0.105	0.171
		Weibull Model			
ŕ	0.027	0.004	<0.001	0.019	0.035
ś	1.427	0.223	<0.001	0.989	1.864
		H_0: ś = 1, *p*-value = 0.056			

Table F3: Estimates of the Effects of Treatments and Demographics
LL = -18366.9

Parameter	Point Estimate	Standard Error	*p*-value	95% Confidence Interval	
δ Constant	0.162	0.050	0.001	0.063	0.260
FED	0.025	0.020	0.216	-0.014	0.064
INFO	-0.045	0.020	0.023	-0.083	-0.006
H_ORDER	-0.041	0.021	0.053	-0.082	0.001
P_HIGH	-0.006	0.020	0.759	-0.046	0.034
RA_FIRST	0.007	0.024	0.768	-0.040	0.055
FEE_HIGH	-0.022	0.023	0.344	-0.068	0.024
FEMALE	-0.043	0.025	0.088	-0.091	0.006
YOUNG	0.003	0.045	0.951	-0.086	0.092
MIDDLE	-0.008	0.036	0.829	-0.079	0.064
OLD	-0.034	0.035	0.329	-0.102	0.034
EDUCATION	-0.011	0.026	0.666	-0.062	0.039
HIGH INCOME	0.006	0.024	0.805	-0.040	0.052
r Constant	0.427	0.166	<0.001	0.102	0.752
RA_FIRST	0.072	0.081	0.378	-0.088	0.232
FEE_HIGH	0.141	0.085	0.097	-0.026	0.307
FEMALE	0.278	0.117	0.018	0.048	0.507
YOUNG	0.015	0.178	0.934	-0.334	0.363
MIDDLE	-0.038	0.126	0.763	-0.286	0.209
OLD	0.070	0.113	0.535	-0.151	0.290
EDUCATION	0.078	0.093	0.400	-0.104	0.259
HIGH INCOME	-0.112	0.081	0.165	-0.270	0.046

Figure F3: Random Coefficient Estimates of Risk Attitudes and Discount Rates

Assuming CRRA utility and Exponential discounting

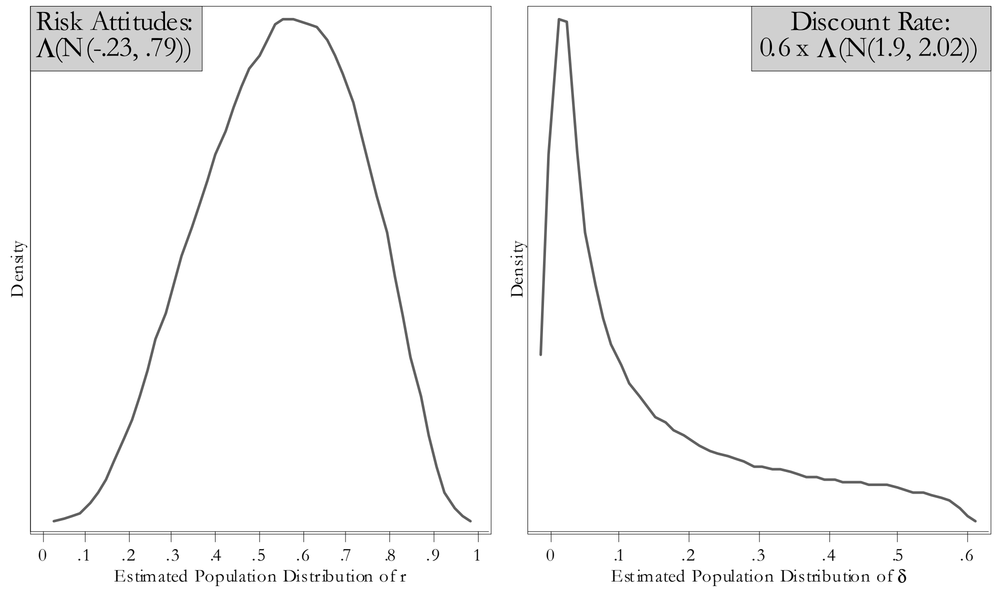

Figure F4: Random Coefficient Estimates of Quasi-Hyperbolic Discounting Model

Assuming CRRA utility and Quasi-Hyperbolic discounting

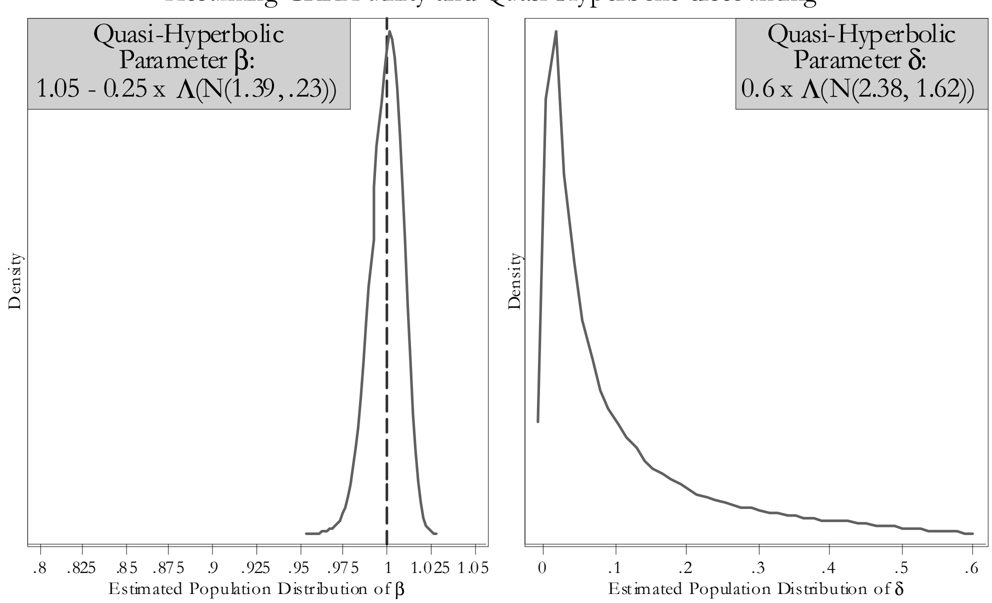

Table F4: ML Estimates of Discounting Models with linear Utility

Parameter	Point Estimate	Standard Error	*p*-value	95% Confidence Interval	
		A. Exponential Discounting (LL = -20814.9; equation (1))			
δ	0.183	0.015	<0.001	0.155	0.212
		B. Quasi-Hyperbolic Discounting (LL = -19705.9; equations (2a) & (2b))			
β	1.005	0.011	<0.001	0.983	1.027
δ	0.150	0.013	<0.001	0.125	0.175
		H_0: β = 1, *p*-value = 0.64			
		C. Fixed Cost Hyperbolic Discounting (LL = -19680.4; equations (3a) & (3b))			
θ	7.086	3.624	0.051	-0.018	14.189
β	0.982	0.034	<0.001	0.916	1.049
δ	0.300	0.123	0.015	0.058	0.542
b	-0.094	0.080	0.243	-0.251	0.064
		H_0: β = 1, *p*-value = 0.60; H_0: β = 1 & b=0, *p*-value = 0.27			
		D. Simple Hyperbolic Discounting (LL = -19705.8; equation (4))			
K	0.180	0.014	<0.001	0.152	0.208
		E. General Hyperbolic Discounting (LL = -19704.6; equation (5))			
α	0.489	0.672	0.467	-0.827	1.805
β	0.200	0.048	<0.001	0.106	0.294
		F. Weibull Discounting (LL = -19707.2; equation (6))			
ŕ	0.166	0.012	<0.001	0.144	0.189
ś	1.053	0.144	<0.001	0.771	1.336
		H_0: ś = 1, *p*-value = 0.71			

Additional References

Andersen, Steffen; Harrison, Glenn W.; Hole, Arne Rise; Lau, Morten I., and Rutström, E. Elisabet, "Non-Linear Mixed Logit," *Theory and Decision*, 73, 2012, 77-96.

Harrison, Glenn W.; Humphrey, Steven J., and Verschoor, Arjan, "Choice Under Uncertainty: Evidence from Ethiopia, India and Uganda," *Economic Journal*, 120, March 2010, 80-104.

Harrison, Glenn W., and Rutström, E. Elisabet, "Expected Utility *And* Prospect Theory: One Wedding and A Decent Funeral," *Experimental Economics*, 12(2), June 2009, 133-158.

Johnson, N.L., "Systems of Frequency Curves Generated by Methods of Translation," *Biometrika*, 36, January 1949, 149-176.

Prelec, Drazen, "The Probability Weighting Function," *Econometrica*, 66, 1998, 497-527.

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