

Eliciting Risk and Time Preferences

by

Steffen Andersen, Glenn W. Harrison, Morten I. Lau and E. Elisabet Rutström[†]

November 2007

Working Paper 05-24, Department of Economics,
College of Business Administration, University of Central Florida, 2005.

Forthcoming, *Econometrica*

Abstract. We design experiments to jointly elicit risk and time preferences for the adult Danish population. Since subjects are generally risk averse, we find that joint elicitation provides estimates of discount rates that are significantly lower than those found in previous studies and more in line with what would be considered as *a priori* reasonable rates. The statistical specification relies on a theoretical framework that involves a latent trade-off between long-run optimization and short-run temptation. Estimation of this specification is undertaken using structural, maximum likelihood methods. Our main results based on exponential discounting are robust to alternative specifications such as hyperbolic discounting. These results have direct implications for attempts to elicit time preferences, as well as debates over the appropriate domain of the utility function when characterizing risk aversion and time consistency.

[†] Department of Economics and Centre for Economic and Business Research, Copenhagen Business School, Copenhagen, Denmark (Andersen); Department of Economics, College of Business Administration, University of Central Florida, USA (Harrison and Rutström) and Department of Economics and Finance, Durham Business School, Durham University, United Kingdom (Lau). Harrison is also affiliated with the Durham Business School, Durham University, UK. E-mail: SA@CEBR.DK, GHARRISON@RESEARCH.BUS.UCF.EDU, M.I.LAU@DURHAM.AC.UK, and ERUTSTROM@BUS.UCF.EDU. Harrison and Rutström thank the U.S. National Science Foundation for research support under grants NSF/IIS 9817518, NSF/HSD 0527675 and NSF/SES 0616746. Harrison and Lau thank the Danish Social Science Research Council for research support under project #24-02-0124. We are grateful to Peter Wakker, Nathaniel Wilcox, six anonymous referees and David Levine for helpful comments. Supporting data and instructions are stored at the ExLab Digital Archive, located at <http://exlab.bus.ucf.edu>.

Utility functions are characterized in three dimensions, reflecting preferences over goods, time and uncertainty. The utility function conventionally characterizes preferences over goods defined by a time period and a state of nature, preferences over the temporal allocation of goods, and preferences over outcomes as realizations of uncertain states of nature. This broad characterization includes most alternatives to conventional expected utility theory.¹ We focus on the utility function for income and background consumption, collapsing the choice over goods down to just one good so that there is no choice option with respect to goods.² We jointly elicit risk and time preferences, using controlled experiments with field subjects in Denmark. Joint elicitation is necessary for identification of both risk and time preference parameters, and we present subjects with one task that identifies their risk attitudes and another task that identifies their discount rate.

We show that *joint elicitation of risk and time preferences results in significantly lower discount rates than separate elicitation of discount rates*. The reason is that we infer discount rates defined in terms of temporally dated utility, rather than assume that individuals are risk neutral and that discount rates are defined in terms of temporally dated income. Since our subjects have concave utility functions, the implied discount rates are lower than when one incorrectly assumes a linear utility function.

We use data from a field methodology developed by Harrison, Lau, Rutström and Sullivan [2005] (HLRS) to elicit both risk and time preferences from the same respondents. They used relatively simple experimental procedures that have evolved in the recent literature to study each. These experimental procedures are presented in section 1, and build on the risk aversion experiments of Holt and Laury [2002] (HL) and the discount rate experiments of Collier and Williams [1999] (CW) and Harrison, Lau and Williams [2002] (HLW). Data is collected in the field in Denmark, to obtain a sample that offers a wider range of individual socio-demographic characteristics than usually found in subject pools recruited in colleges, as well as a sample that can

¹ The major exception for present purposes is the approach of Kreps and Porteus [1978], which allows for preferences over the timing of the resolution of uncertainty.

² Background consumption is the optimized consumption stream based on wealth and income that is perfectly anticipated, before allowing for the effects of the money offered in the experimental tasks.

be used to make inferences about the preferences of the adult population of Denmark. These experiments are “artefactual field experiments” in the terminology of Harrison and List [2004], since we essentially take lab experiments to field subjects.

Our statistical specification uses a theoretical framework that involves a trade-off between long-run optimization and short-run temptation, following theories presented by Fudenberg and Levine [2006] (FL), and Benhabib and Bisin [2005]. We view the risk responses as being made with a strong short-run focus since subjects are paid for these responses at the end of the experiment, implying the presence of immediate temptation. The income earned from the risk task is therefore integrated with short-run background consumption. On the other hand, the discount rate responses have payoff consequences at least one month into the future which removes immediate temptation, and income from these responses are therefore integrated with background consumption over a longer period of time. The relationship between risk and time preferences is initially specified by an exponential discount function and an explicit utility function. We evaluate the implied latent theoretical specification in section 2 using structural maximum likelihood methods. These estimates show moderate risk aversion and an average discount rate of 10.1% per annum. This estimate is significantly lower than the average discount rate of 25.2% per annum inferred from similar procedures that assume risk neutrality.

Our estimates rely on parametric functional forms, which is just to say that some theoretical structure is needed to measure these latent preferences correctly. So it is important to consider alternative specifications and functional forms, even if the ones we initially chose are standard. In section 3 we consider models that (i) allow for observable individual characteristics, (ii) reflect unobservable differences in preferences, (iii) use hyperbolic specifications of the discount rate function, and finally (iv) allow for more than one latent data-generating process. We find that our main result is robust to the use of these alternative specifications: discount rates defined over utility streams are significantly lower than discount rates defined over income streams.

There are only a few studies that address the joint elicitation of risk and time preferences

directly using monetary incentives and procedures familiar to experimental economists. None of these studies consider the formal theoretical link between elicited risk attitudes and individual discount rates that is our focus. We review this related literature in section 4. Finally, section 5 concludes with implications of our research.

1. Experimental Procedures for Eliciting Risk and Time Preferences

Our experimental procedures are documented in detail in HLRS, so we focus here just on the basics. In brief, each subject was asked to respond to four risk aversion tasks and six discount rate tasks. Each such task involved a series of binary choices, typically 10 per task. Thus each subject typically provides 100 binary choices that can be used to infer risk and time preferences.

A. Risk Preferences: Measuring Risk Aversion

Holt and Laury [2002] devise a simple experimental measure for risk aversion using a multiple price list (MPL) design.³ Each subject is presented with a choice between two lotteries, which we can call A or B. Table 1 illustrates the basic payoff matrix presented to subjects in our experiments.⁴ The first row shows that lottery A offered a 10% chance of receiving 2,000 DKK and a 90% chance of receiving 1,600 DKK. The expected value of this lottery, EV^A , is shown in the third-last column as 1,640 DKK, although the EV columns were not presented to subjects. Similarly, lottery B in the first row has chances of payoffs of 3,850 and 100 DKK, for an expected value of 475 DKK. Thus the two lotteries have a relatively large difference in expected values, in this case 1,165 DKK. As one proceeds down the matrix, the expected value of both lotteries increases,

³ The MPL appears to have been first used in pricing experiments by Kahneman, Knetsch and Thaler [1990], and has been adopted in recent discount rate experiments by CW. It has a longer history in the elicitation of hypothetical valuation responses in “contingent valuation” survey settings, as discussed by Mitchell and Carson [1989; p. 100, fn. 14]. Andersen, Harrison, Lau and Rutström [2006] examine the properties of the MPL procedure in detail, and Harrison and Rutström [2007] review alternative elicitation procedures for risk attitudes.

⁴ As explained in HLRS the task also varied across subjects in terms of the width of the intervals and the number of rows.

but the expected value of lottery B becomes greater relative to the expected value of lottery A.

The subject chooses A or B in each row, and one row is later selected at random for payout for that subject. The logic behind this test for risk aversion is that only risk-loving subjects would take lottery B in the first row, and only very risk-averse subjects would take lottery A in the second last row. Arguably, the last row is simply a test that the subject understood the instructions, and has no relevance for risk aversion at all. A risk neutral subject should switch from choosing A to B when the EV of each is about the same, so a risk-neutral subject would choose A for the first four rows and B thereafter.

We take each of the binary choices of the subject as the data, and estimate the parameters of a latent utility function that explains those choices using an appropriate error structure to account for the panel nature of the data. Once the utility function is defined, for a candidate value of the parameters of that function, we can construct the expected utility of the two gambles, and then use a linking function to infer the likelihood of the observed choice. We discuss statistical specifications in more detail in section 2.

We undertake four separate risk aversion tasks with each subject, each with different prizes designed so that all 16 prizes span the range of income over which we seek to estimate risk aversion. The four sets of prizes are as follows, with the two prizes for lottery A listed first and the two prizes for lottery B listed next: (A1: 2000 DKK, 1600 DKK; B1: 3850 DKK, 100 DKK), (A2: 2250 DKK, 1500 DKK; B2: 4000 DKK, 500 DKK), (A3: 2000 DKK, 1750 DKK; B3: 4000 DKK, 150 DKK), and (A4: 2500 DKK, 1000 DKK; B4: 4500 DKK, 50 DKK). At the time of the experiments, the exchange rate was approximately 6.55 DKK per U.S. dollar, so these prizes range from approximately \$7.65 to \$687.

We ask the subject to respond to all four risk aversion tasks and then randomly decide which task and row to play out. In addition, the large incentives and budget constraints precluded paying all subjects, so each subject is given a 10% chance to actually receive the payment associated with his

decision.⁵

B. Time Preferences: Measuring Individual Discount Rates

The basic experimental design for eliciting individual discount rates (IDRs) was introduced in CW and expanded in HLW. Subjects in our experiments were given payoff tables such as the one illustrated in Table 2, with 10 symmetric intervals. In this example, Option A offered 3000 DKK in one month and Option B paid 3000 DKK + x DKK in seven months, where x ranged from annual rates of return of 5% to 50% on the principal of 3000 DKK, compounded quarterly to be consistent with general Danish banking practices on overdraft accounts. The payoff tables provided the annual and annual effective interest rates for each payment option, and the experimental instructions defined these terms by way of example.⁶ Subjects were asked to choose between Option A and B for each of the 10 payoff alternatives, and one decision row was selected at random to be paid out at the chosen date. If a risk-neutral subject prefers the 3000 DKK in one month then, using the standard model of the previous experimental literature, we can infer that the annual discount rate is higher than $x\%$; otherwise, we can infer that it is $x\%$ or less.⁷ We examine the theoretical basis for these direct inferences about discount rates in section 2.

We use the multiple-horizon treatment from HLW. From the perspective of the task faced by the subjects, the only variations are that the instrument is now computerized, and subjects are presented with 6 discount rate tasks, corresponding to 6 different time horizons: 1 month, 4 months,

⁵ There is considerable behavioral evidence that rewarding subjects by selecting one task at random for payment does not distort choices, even though it does make the overall experiment a compound lottery. See Harrison, Lau and Rutström [2007; fn.16] for evidence on this issue for the risk aversion instrument we used here, and Harrison and Rutström [2007; §2.6] for similar evidence in comparable lottery choice tasks.

⁶ CW and HLW provided annual and annual effective interest rates to help subjects compare lab and field investments. This feature may reduce comparison errors and CW find that providing information on interest rates has a significant negative effect on elicited discount rates.

⁷ We assume that the subject does not have access to perfect capital markets, as explained in CW (p.110) and HLW (p.1607ff.). This assumption is plausible, but also subject to checks from responses to the financial questionnaire that CW, HLW and we ask each subject to complete. The effects of allowing for field borrowing and lending opportunities on elicited discount rates for risk neutral subjects are discussed by CW and HLW; Harrison, Harstad and Rutström [2004] discuss the general implications of allowing for extra-experimental trading opportunities on inferences from experimental responses.

6 months, 12 months, 18 months, and 24 months.⁸ In each task subjects are provided two future income options rather than one “instant income” option and one future income option. We follow HLW and use a delay of one month to the early income option in all tasks. For example, they were offered 3000 DKK in one month and 3000 DKK + x DKK in 7 months, so that we interpret the revealed discount rate as applying to a time horizon of 6 months. This avoids the potential problem of the subject facing extra risk or transactions costs with the future income option, as compared to the “instant” income option.⁹ If the delayed option were to involve such additional transactions costs, then the revealed discount rate would include these subjective transactions costs. By having both options presented as future income we hold these transactions costs constant.

Each subject responded to all six discount rate tasks and one task and row was chosen at random to be played out. Future payments to subjects were guaranteed by the Danish Ministry of Economic and Business Affairs, and made by automatic transfer from the Ministry’s bank account to the subject’s bank account. This payment procedure is similar to a post-dated check, and automatic transfers between bank accounts are a common procedure in Denmark. Finally, each subject was given a 10% chance to receive actual payment. Thus, each subject faced a 10% chance of receiving payment in the risk preference task as well as a 10% chance in the time preference task.

Our estimation strategy is the same as for the lottery task. We take each of the binary choices of the subject as data, and estimate the parameters with an error structure that recognizes the panel nature of the data. Risk attitudes and discount rates are estimated jointly. In effect, the lottery tasks identify risk attitudes and the intertemporal tasks identify discount rates.

⁸ The design mimics the format used by HL in their risk aversion experiments: in that case the rows reflected different probabilities of each prize, and in this case the rows reflect different annual effective rates of return. We exploit this similarity of format in the design of our computerized interface to subjects, and in the use of trainers in the risk aversion task as a generic substitute for trainers in the discount rate task.

⁹ These transactions costs are discussed in CW, and they include simple things such as remembering to pick up the delayed payment as well as more complex things such as the credibility of the money actually being paid in the future. As discussed in HLRS the design of our experiment was intended to make sure that the credibility of receiving the money in the future was high. These considerations may be important in a field context, particularly in less developed countries.

C. Experimental Procedures and Data

The sample for the field experiments was designed to generate a representative sample of the adult Danish population between 19 and 75 years of age. The experiments were conducted over 20 sessions, between June 2 and June 24, 2003, in 19 locations spread over Denmark, and a total of 253 subjects participated. Details of the procedures, instructions, and sample frame are provided in Supplementary Materials available online.

2. Identifying Risk and Time Preferences

A. General Statement

Consider the identification of risk and time preferences in the canonical case of mainstream economic theory. Specifically, if we assume that Expected Utility Theory (EUT) holds for the choices over risky alternatives and that discounting is exponential then the subject is indifferent between two income options M_t and $M_{t+\tau}$ if and only if

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

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 (1) 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 (1) 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 .

Most analyses of discounting models implicitly assume that the individual is risk neutral,¹⁰ so

¹⁰ See Keller and Strazzera [2002; p. 148] and Frederick, Loewenstein and O'Donoghue [2002;]

that (1) is instead written in the more familiar form

$$M_t = (1/(1+\delta)^\tau) M_{t+\tau} \quad (2)$$

where δ is the discount rate that makes the present value of the two monetary outcomes M_t and $M_{t+\tau}$ equal.

To state the obvious, (1) and (2) are not the same. As one relaxes the assumption that the decision maker is risk neutral, it is apparent from Jensen's Inequality that the implied discount rate decreases if $U(M)$ is concave in M . Thus one cannot infer the level of the individual discount rate without knowing or assuming something about their risk attitudes. This identification problem implies that risk attitudes and discount rates cannot be estimated based on discount rate experiments alone, but separate tasks to identify the influence of risk preferences must also be implemented.

B. Parametric Structure

We can quickly put some familiar parametric structure on this statement of the identification problem. Let the utility function be the constant relative risk aversion (CRRA) specification

$$U(M) = (\omega + M)^{1-r} / (1-r) \quad (3)$$

for $r \neq 1$, where r is the CRRA coefficient.¹¹ Background consumption, ω , is assumed to be zero in some studies and to represent lifetime wealth in other studies. We assume that $\omega > 0$ and view this specification as a "reduced form" model that is consistent with a variety of structural models, as discussed in section C. With this functional form, $r=0$ denotes risk neutral behavior, $r>0$ denotes

p.381ff.] for an explicit statement of this assumption, which is often implicit in applied work, as well as examples of the wide use of the exponential discounting function. We refer to risk aversion and concavity of the utility function interchangeably, but it is concavity that is central (the two can differ for non-EUT specifications).

¹¹ There are numerous applications of the CRRA specification under Expected Utility Theory and Prospect Theory. Gollier [2001; p.27] notes the different asymptotic properties of this CRRA specification when r is higher and lower than one. When $r < 1$ utility goes from 0 to ∞ as income goes from 0 to ∞ , and when $r > 1$ utility goes from *minus* ∞ to 0 as income goes from 0 to ∞ . An alternative CRRA function is $U(M) = (\omega + M)^\rho$, where $\rho \neq 0$ and $(1-\rho)$ is the CRRA coefficient. Wakker [2006] provides a valuable exposition of the formal properties of this venerable function, and some variants found in the literature. While the latter function has the unfortunate property that marginal utility is negative for $\rho < 0$, this is not the case for the utility function (3) that we use. In addition, even though the CRRA specification (3) can generate negative discount rates for extreme risk attitudes, such as $r > 1$, in the absence of background consumption ω , it is well-behaved when background consumption is included, as it is here.

risk aversion, and $r < 0$ denotes risk loving behavior.

To relate this specification to the risk aversion choices in our experiment, one can calculate the implied bounds on the CRRA coefficient for each row in Table 1, and these are in fact reported by HL [2002; Table 3]. These intervals are shown in the final column of Table 1, assuming that background consumption ω is zero. Thus, for example, a subject that made 5 safe choices and then switched to the risky alternatives would have revealed a CRRA interval between 0.14 and 0.41, and a subject that made 7 safe choices would have revealed a CRRA interval between 0.68 and 0.97, and so on. Thus the binary choices of the subject can be explained by different values of the CRRA coefficient, and the coefficient estimated using standard maximum likelihood procedures (explained in detail below). For positive background consumption levels the same observed choices would imply higher absolute values of the CRRA coefficients, i.e. more curvature of the utility functions.

There is evidence from the lab and field that subjects are risk averse over stakes ranging between pennies and several hundred dollars. Holt and Laury [2002][2005] produced the most widely cited evidence from the lab, and they show that subjects are moderately risk averse. Harrison, Lau and Rutström [2007] find comparable results with artefactual field experiments conducted with the adult Danish population used for our analysis. This literature also offers some evidence of lower estimates of relative risk aversion when the stakes in the experimental task are reduced significantly, which may cause one to question our use of the restrictive CRRA function. However, Harrison, Lau and Rutström [2007] find that CRRA holds locally over the domain of stakes offered here, and we adopt this specification because it is convenient and popular in the theoretical and empirical literature on risk preferences. We do not claim that the particular estimates provided here would hold globally if the stakes in our tasks are reduced or increased by a significant amount.

C. Theoretical Issues

The theoretical problem with this empirical evidence is that one expects decision-makers to be essentially risk neutral over small stakes such as these if they integrate the lottery prizes with

lifetime wealth (or consumption over any period exceeding several months). If the decision-maker is not risk neutral with respect to small stake lotteries, one can construct absurd predictions that are counter-intuitive (Rabin [2000]).

Turning to time preferences, there is evidence that decision makers exhibit a “passion for the present” when offered choices between monetary amounts today or in the future, but that they do not exhibit such a passion when offered choices between monetary amounts at different dates in the future. That is, the discount rate required to rationalize the choice of money *today* or in the future is extremely high (in the order of hundreds of percent p.a., or even thousands of percent), but the discount rate required to rationalize the choice of money at *two distinct future dates* is relatively low and constant with respect to the implied time delay (Frederick, Loewenstein and O’Donoghue [2002] and HLW). This sharp drop in discount rates is most popularly characterized by “quasi-hyperbolic preference” structures. A time delay on the early payment is referred to as a “front end delay,” so these results hinge on whether that front end delay is zero or positive.

Fudenberg and Levine [2006] propose a dual-self model¹² that addresses these two “elephants in the room.” They propose a framework in which a decision maker does not always make decisions in a purely myopic manner, nor does he always optimize over lifetime wealth. They acknowledge that each person has an interest in long term goals, but is at the same time tempted by more immediate consumption opportunities. One of the implications of the FL model is that income received in an experimental setting is completely integrated with lifetime wealth as an argument of utility only when short term temptation is completely absent or under self-control. On the other hand, in some situations the short term temptation completely dominates, and there is no consideration for long term goals, implying that income is not integrated into lifetime wealth. The

¹² The concept of “dual selves” has a long lineage in behavioral economics, and findings from neuroscience suggest that multiple brain systems interact when subjects make economic decisions (Cohen [2005]). An alternative interpretation of the concept “dual selves” would be a single decision maker that has dual cognitive processes that are activated under different conditions. This interpretation is consistent with the literature on dual process theories of mind in psychology and economics (e.g., Barrett, Tugade and Engle [2004] and Benhabib and Bisin [2005]).

costliness of exercising self-control determines the trade-off between lifetime wealth maximization and short term temptation, so final decisions tend neither to reflect pure short term temptation nor pure lifetime wealth maximization.

Our empirical model reflects such tradeoffs by varying how income from the experimental task is integrated into other extra-experimental income and wealth. We are agnostic with respect to the mechanism by which the integration takes place, recognizing only that when money is paid out immediately the shorter term temptations are stronger than when money is paid out with some temporal delay. Subjects in our experiments are confronted with lottery choices for which they are paid immediately, as well as non-stochastic asset choices for which they are paid with at least a one month delay. We specify one utility function for the risk aversion task and another utility function, with a less binding liquidity constraint, for the discount rate task. This distinction is justified by the temporal difference in the receipt of the payoffs in the two sets of decisions. In other settings one could also imagine the distinction being justified by the size of the stakes in the payoffs, although that is not a factor in our experiments.

We then estimate the parameters of these utility functions simultaneously. The parameters to be estimated include the risk parameter r and the discount rate δ . Immediate experimental income is integrated directly with background consumption ω , which is an exogenous parameter in the model. In the utility function for the risk aversion task, immediate experimental income can be thought of as divided evenly across η periods of time, and the fraction $1/\eta$ of this income is integrated with background consumption ω . We set $\eta=1$ throughout, but retain the notation to contrast with the treatment of delayed experimental income from the discount rates tasks. Delayed experimental income is divided evenly across λ periods of time, and the fraction $1/\lambda$ of this income is integrated with background consumption ω . We typically consider cases in which $\lambda>1$, since the dual self modeling assumption requires that $\lambda>\eta$. We parametrically vary this factor λ to trace the effect it has on the estimated risk parameter and discount rate.¹³ Thus, the utility functions differ *only* in the

¹³ An important maintained assumption in our empirical model is that the risk parameter r is constant

treatment of how background consumption ω is integrated with experimental income from the different tasks.

It is important to be precise on how this theoretical framework is applied in experimental settings, since considerable confusion has arisen about the theoretical interpretation of risk and time preferences elicited in experiments. If an agent who is maximizing utility subject to a lifetime budget constraint is given the lotteries that are typical in lab experiments he should be risk-neutral, since the stakes would be spread over the expected lifetime and amount to minuscule amounts per period. In FL this type of agent is modeled as the long-run self; an alternative strategy in Benhabib and Bisin [2005] (BB) models the agent as activating a controlled cognitive process which is immune to temptation. If this agent is also asked to make choices between monetary amounts to be paid at different time periods, he would exhibit constant discounting, whether or not a front end delay was used.

Conversely, if an agent who is completely open to temptation is asked to make the lottery choices in the risk aversion task, where prizes are to be resolved today, he would not necessarily be risk neutral. And if he is asked to make choices over monetary amounts at different time horizons, he would pick the amount with no front end delay, and would be indifferent if offered choices between amounts at two future dates. This agent, completely under the spell of temptation, is modeled as the short-term decision maker in FL or as an agent relying on an automatic cognitive process susceptible to temptation in BB.

So clearly neither FL's long-run self nor the short-run self would be able to explain all of the experimental choices we observe. Instead, we observe what appears to be the outcome of a decision

over time and planning horizons. Thus we assume the same r applies to utility over the lottery payoffs, which are paid immediately, as to utility over the asset choices, which all have a front-end delay. This assumption is plausible from both an empirical and a theoretical perspective. Andersen, Harrison, Lau and Rutström [2008] use data from a panel experiment with similar monetary incentives and find some variation in risk attitudes over time, but they do not detect a general tendency for risk attitudes to increase or decrease over a 17-months time span. Theoretically, it allows us to focus on the role of differences in consumption smoothing between the two selves of our agent. Rich as our experimental design is, we do not believe it would be possible to identify differences in the risk parameter as well as differences in consumption smoothing.

process where temptation and long-run considerations are simultaneously involved, such as when the choices of FL's short-run self are imperfectly constrained by a *latent* long-run self through some self-control mechanism. Exactly how this self-control is effected is left vague, since one could imagine a wide range of cognitive or logistical processes that might do the job. But if the cost of self-control is positive, temptation will continue to influence decisions but possibly tempered by considerations of the longer-run.

One attractive feature of this theoretical specification is that the usual assumptions of the experimental literature emerge when we set $\eta=\lambda=1$. It allows us to cleanly evaluate the implications of taking these theoretical issues seriously when making inferences from experimental data. Since $\eta=1$ throughout, this special case is obtained later when we report estimates conditional on assuming that $\lambda=1$.

D. Statistical Specification

We apply the dual-self model in our analysis and explicitly recognize the different trade-offs between short-run temptation and long-run goals that agents make across the different tasks. The risk aversion responses are those of an unconstrained short-run self falling prey to temptation. The payment is immediate and may therefore be seen as integrated with whatever pocket-cash the subject normally has access to.¹⁴ The discount rate tasks all used a front end delay, making immediate temptation irrelevant. These delayed payments are therefore likely to be integrated with background consumption over a longer period of time, which may fall well short of the remaining lifetime if there are spending goals in intermediate time periods that exert some temptation.

We can write out the likelihood function for the choices that our subjects made and jointly estimate the risk parameter r and the discount rate δ . Consider first the contribution to the overall likelihood from the risk aversion responses. Probabilities for each outcome M_i , $p(M_i)$, are those that

¹⁴ It may be the case that the exertion of self-control is motivated by the stakes offered so that with larger stakes and greater self-control the relevant background consumption may reflect longer-run goals even when the payment is immediate. We do not model this possibility here, however.

are induced by the experimenter, so expected utility is simply the probability weighted utility of each outcome in each lottery. Since there were two outcomes in each lottery, the EU for lottery i is

$$EU_i = \sum_{j=1,2} [p(M_j) \times U(\omega+M_j/\eta)] = \sum_{j=1,2} [p(M_j) \times U(\omega+M_j)]. \quad (4)$$

We drop the symbol η for simplicity, since $\eta=1$ and is only needed to remind us that $\lambda \geq \eta$.

A simple stochastic specification from Holt and Laury [2002] is used to specify likelihoods conditional on the model. The EU for each lottery pair is calculated for candidate estimate of r and ω , and the ratio

$$\nabla EU = EU_B^{1/\mu} / (EU_A^{1/\mu} + EU_B^{1/\mu}) \quad (5)$$

calculated, where EU_A refers to Option A and EU_B refers to Option B, and μ is a structural “noise parameter” used to allow some errors from the perspective of the deterministic EUT model. The index ∇EU is in the form of a cumulative probability distribution function defined over differences in the EU of the two lotteries and the noise parameter μ .¹⁵ Thus, as $\mu \rightarrow 0$ 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. This is one of several different types of error story that could be used.¹⁶ The index in (5) is linked to observed choices by specifying that the Option B lottery is chosen when $\nabla EU > 1/2$.

Thus the likelihood of the risk aversion responses, conditional on the EUT and CRRA specifications being true, depends on the estimates of r and μ , and the observed choices. If we ignore the responses that reflect indifference, the conditional log-likelihood is

$$\ln L^{RA}(r, \mu; y, \omega, \mathbf{X}) = \sum_i [(\ln(\nabla EU) \mid y_i = 1) + (\ln(1 - \nabla EU) \mid y_i = -1)] \quad (6)$$

¹⁵ An alternative approach might be to define an index as the EU difference, and then specify some cumulative distribution function to link it to the observed choices. For example, the cumulative standard normal distribution leads to the probit specification.

¹⁶ See Harless and Camerer [1994], Hey and Orme [1994] and Loomes and Sugden [1995] for the first wave of empirical studies including some formal stochastic specification in the version of EUT tested. There are several species of “errors” in use, reviewed by Hey [1995][2002], Loomes and Sugden [1995], Ballinger and Wilcox [1997], and Loomes, Moffatt and Sugden [2002]. Some place the error at the final choice between one lottery or the other after the subject has decided deterministically 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.

where $y_i = 1(-1)$ denotes the choice of the Option B (A) lottery in risk aversion task i , and \mathbf{X} is a vector of individual characteristics. The subjects were told at the outset that any expression of indifference would mean that if that choice was selected to be played out the experimenter would toss a fair coin to make the decision for them. Hence one can modify the likelihood to take these responses into account by recognizing that such choices implied a 50:50 mixture of the likelihood of choosing either lottery

$$\ln L^{RA}(\mathbf{r}, \boldsymbol{\mu}; \mathbf{y}, \boldsymbol{\omega}, \mathbf{X}) = \sum_i [(\ln (\nabla EU) \mid y_i = 1) + (\ln (1 - \nabla EU) \mid y_i = -1) + (\frac{1}{2} \ln (\nabla EU) + \frac{1}{2} \ln (1 - \nabla EU) \mid y_i = 0)] \quad (6')$$

where $y_i = 0$ denotes the choice of indifference. Only 4.6% of the observed choices in our experiments were expressions of indifference, but it is appropriate to use (6') to account for them.¹⁷

A similar specification is employed for the discount rate choices. Equation (4) is replaced by the discounted utility of each of the two options, conditional on some assumed discount rate, and equation (5) is defined in terms of those discounted utilities instead of the expected utilities. If we assume that the subject integrates the two monetary amounts in the discount rate tasks with background consumption $\boldsymbol{\omega}$ over λ periods of time, then Option A that is paid out at time t provides the following intertemporal sequence of utility

$$[U_t(\boldsymbol{\omega} + M_A/\lambda), U_{t+1}(\boldsymbol{\omega} + M_A/\lambda), \dots, U_{t+\lambda-1}(\boldsymbol{\omega} + M_A/\lambda)] [U_{t+\tau}(\boldsymbol{\omega}), U_{t+\tau+1}(\boldsymbol{\omega}), \dots, U_{t+\tau+\lambda-1}(\boldsymbol{\omega})]$$

and Option B that is paid out at time $t+\tau$ provides the intertemporal sequence of utility

$$[U_t(\boldsymbol{\omega}), U_{t+1}(\boldsymbol{\omega}), \dots, U_{t+\lambda-1}(\boldsymbol{\omega})] [U_{t+\tau}(\boldsymbol{\omega} + M_B/\lambda), U_{t+\tau+1}(\boldsymbol{\omega} + M_B/\lambda), \dots, U_{t+\tau+\lambda-1}(\boldsymbol{\omega} + M_B/\lambda)]$$

where M_A and M_B are the monetary amounts in the choice tasks presented to subjects, illustrated in Table 2.¹⁸ The discounted utility of Option A is then given by

¹⁷ Our treatment of indifferent responses uses the specification developed by Papke and Wooldridge [1996; equation 5, p.621] for fractional dependant variables. Alternatively, one could follow Hey and Orme [1994; p.1302] and introduce a new parameter τ to capture the idea that certain subjects state indifference when the latent index showing how much they prefer one lottery over another falls below some threshold τ in absolute value. This is a natural assumption to make, particularly for the experiments they ran in which the subjects were told that expressions of indifference would be resolved by the experimenter, but not told how the experimenter would do that (p.1295, footnote 4). It adds one more parameter to estimate, but for good cause.

¹⁸ The even division of income over the λ periods of time is a convenient assumption that simplifies

$$PV_A = \sum_{i=\{t, \dots, t+\lambda-1\}} [(1/(1+\delta))^{(i-t)} U(\omega+M_A/\lambda)] + \sum_{i=\{t+\tau, \dots, t+\tau+\lambda-1\}} [(1/(1+\delta))^{(i-t)} U(\omega)] \quad (7)$$

and the discounted utility of Option B is

$$PV_B = \sum_{i=\{t, \dots, t+\lambda-1\}} [(1/(1+\delta))^{(i-t)} U(\omega)] + \sum_{i=\{t+\tau, \dots, t+\tau+\lambda-1\}} [(1/(1+\delta))^{(i-t)} U(\omega+M_B/\lambda)] \quad (8)$$

where the utility function is assumed to be stationary over time.¹⁹

The parameter λ defines the number of periods over which the two *delayed* monetary amounts in the discount rate choices are integrated with background consumption ω . If $\lambda=1$ then the period of assumed consumption is the same across the risk aversion and discount rate tasks, since the period of assumed consumption in the risk aversion task, η , is equal to 1 throughout. But we generally expect λ to be greater than 1 due to the differences in the delay with which the prizes are paid out: no delay for the risk aversion task, and at least 30 days for the discount rate task. For very large values of λ we approach lifetime wealth. One interpretation of η and λ is that they represent the time horizon over which the subject is optimizing in each task. For example, if the background consumption in the risk aversion task corresponds to the amount of money a subject would spend in a day, then $\eta=1$ by definition, and λ would correspond to the number of days over which subjects expect to spend the earnings in the discount rate task. This specification is different from the common assumption in the discounting literature that subjects are completely liquidity constrained and consume the monetary amounts in options A and B at the time stated in the

the mathematical expressions. We could also divide income unevenly over the λ periods of time. Indeed, the special case where subjects spend the income when it is paid out takes us back to the EUT specification where $\lambda=1$. Our approach to consumption smoothing should be viewed as the simplest possible way in which one could add structure to the experimental tasks we employ. The objective is to be able to address the connection between choices made in the experimental tasks defined over monetary payments, and a latent economic structure in which utility is defined over consumption streams. We noted earlier our assumption that subjects do not have access to perfect capital markets. The special case $\eta=\lambda=1$, common in the experimental literature, in effect views monetary payments as consumption streams. Many find this identifying assumption unattractive theoretically. Our approach can be viewed as one tractable way to explore the sensitivity of inferences about risk attitudes and discount rates to relaxations of that assumption. In fact, there is virtually no data on the time-path of consumption flows for individuals or households at a detailed level. Time-use surveys are focused on work and non-work activities, and not expenditure levels (e.g., Hamermesh and Pfann [2005]). Specialized data sets do provide some longitudinal consumption and income data, such as those on food consumption utilized by Shapiro [2005] to investigate discount rates defined over daily caloric intake, but they are obviously limited in scope.

¹⁹ To be precise, due to the inclusion of a front-end delay both equation (7) and (8) should include an additional discounting factor back to the time when the decision is made. The additional discounting factor cancels out in exponential discounting models, so for ease of exposition this is not shown here.

instrument, and do not smooth consumption over time (Frederick, Loewenstein and O'Donoghue [2002; p. 380]).²⁰ That assumption, again, is just the special case in which $\eta=\lambda=1$.

The parametric form for the utility function in (7) and (8) is the CRRA form given in (3)²¹, so we can rewrite these as

$$PV_A = \sum_{i=\{t, \dots, t+\lambda-1\}} [(1/(1+\delta))^{(i-t)} (\omega+M_A/\lambda)^{(1-r)}] + \sum_{i=\{t+\tau, \dots, t+\tau+\lambda-1\}} [(1/(1+\delta))^{(i-t)} \omega^{(1-r)}] \quad (7')$$

$$PV_B = \sum_{i=\{t, \dots, t+\lambda-1\}} [(1/(1+\delta))^{(i-t)} \omega^{(1-r)}] + \sum_{i=\{t+\tau, \dots, t+\tau+\lambda-1\}} [(1/(1+\delta))^{(i-t)} (\omega+M_B/\lambda)^{(1-r)}]. \quad (8')$$

It is useful to take a closer look at the relation between δ and λ in this specification. The subject is indifferent between the two amounts M_A and M_B if and only if (7')=(8') which can be written as

$$(\omega+M_A/\lambda)^{(1-r)} + (1/(1+\delta)^\tau) \omega^{(1-r)} = \omega^{(1-r)} + (1/(1+\delta)^\tau) (\omega+M_B/\lambda)^{(1-r)},$$

and the discount factor can then be expressed as a function of the two amounts M_A and M_B and the time period λ

$$(1/(1+\delta)^\tau) = [\omega^{(1-r)} - (\omega+M_A/\lambda)^{(1-r)}] / [\omega^{(1-r)} - (\omega+M_B/\lambda)^{(1-r)}].$$

It is easy to show that the discount factor in the risk neutral case ($r=0$) is equal to M_A/M_B .

L'Hôpital's rule can be used to compute the limit value of the discount factor as λ goes to infinity.

Let $f(\lambda)$ denote the numerator and $g(\lambda)$ denote the denominator. We see that $f(\infty)=g(\infty)=0$, and

L'Hôpital's rule then states that

$$\lim_{\{\lambda \rightarrow \infty\}} f(\lambda)/g(\lambda) = f'(\infty)/g'(\infty).$$

Differentiating the denominator and numerator with respect to λ gives

$$f'(\lambda)/g'(\lambda) = (M_A/M_B) [(\omega+M_A/\lambda)/(\omega+M_B/\lambda)]^{(1-r)} \Rightarrow f'(\infty)/g'(\infty) = M_A/M_B.$$

Hence, independent of the risk preferences, the discount factor approaches the value for the risk neutral case (M_A/M_B) when the monetary outcomes in the discount rate tasks are integrated with background consumption over an infinite horizon. For finite values of λ one can show that

²⁰ Although we relax the assumption that subjects spend additional income at the time it is paid out, we are not free to relax the assumption of binding liquidity constraints entirely. Our estimates are restricted by the assumption that subjects are liquidity constrained beyond the time periods η and λ .

²¹ One can theoretically apply different CRRA specifications across the dual selves, but our statistical approach restricts us to use a single CRRA specification for both selves because the risk aversion tasks are used to identify the r -parameter and the discount rate tasks are used to identify the δ -parameter.

$f'(\lambda)/g'(\lambda) > M_A/M_B$ when $r > 0$ and $f'(\lambda)/g'(\lambda) < M_A/M_B$ when $r < 0$, i.e. the discount factor is higher than the risk neutral level if the utility function is concave ($r > 0$) and vice versa when the utility function is convex ($r < 0$).

An index of the difference between these present values, conditional on r and δ , can then be defined as

$$\nabla PV = PV_B^{1/\nu} / (PV_A^{1/\nu} + PV_B^{1/\nu}) \quad (9)$$

where ν is a noise parameter for the discount rate choices, just as μ was a noise parameter for the risk aversion choices. It is not obvious that $\mu = \nu$, 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 $\mu > \nu$. 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.

Thus the likelihood of the discount rate responses, conditional on the EUT, CRRA and exponential discounting specifications being true, depend on the estimates of r , δ , μ and ν , given the assumed value of ω , λ and the observed choices. If we ignore the responses that reflect indifference, the conditional log-likelihood is

$$\ln L^{DR}(r, \delta, \mu, \nu; y, \omega, \lambda, \mathbf{X}) = \sum_i [(\ln(\nabla PV) | y_i = 1) + (\ln(1 - \nabla PV) | y_i = -1)] \quad (10)$$

where $y_i = 1(-1)$ again denotes the choice of Option B (A) in discount rate task i , and \mathbf{X} is a vector of individual characteristics. We can easily add responses that reflect indifference to the log-likelihood function and get

$$\ln L^{DR}(r, \delta, \mu, \nu; y, \omega, \lambda, \mathbf{X}) = \sum_i [(\ln(\nabla PV) | y_i = 1) + (\ln(1 - \nabla PV) | y_i = -1) + (\frac{1}{2} \ln(\nabla PV) + \frac{1}{2} \ln(1 - \nabla PV) | y_i = 0)] \quad (10')$$

where $y_i = 0$ denotes the choice of indifference.

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

$$\ln L(r, \delta, \mu, \nu; y, \omega, \lambda, \mathbf{X}) = \ln L^{RA} + \ln L^{DR} \quad (11)$$

and maximized using standard numerical methods. Our implementation uses version 10 of *Stata*.²² The statistical specification allows for the possibility of correlation between responses by the same subject.²³

E. Estimates

We assume that income from the risk aversion tasks is spent in one day, i.e. $\eta=1$ by definition and $\lambda \geq 1$ corresponds to the number of days over which income from the discount rate tasks is spent. Fudenberg and Levine [2007] consider a similar time horizon for the short-run self, and this assumption seems reasonable given the stakes in the experimental tasks and our payment methods. Using data from the household expenditure survey at Statistics Denmark we find that per capita consumption of private non-durable goods on an average daily basis was equal to 118 DKK in 2003. We use this value of daily background consumption ω in our estimations. We assume that $\lambda=1$ in the baseline, which means that delayed income from the discount rate tasks is integrated with daily background consumption over one day. In addition to our baseline estimation, we also consider values of $\lambda > 1$ and variations in ω , as well as a restricted specification assuming risk neutrality.

Table 3 presents maximum likelihood estimates from our experiments.²⁴ Panel A presents

²² We document the *Stata* syntax for this estimation at <http://exlab.bus.ucf.edu>. We also provide all source code and data for the estimates reported here.

²³ The use of clustering to allow for “panel effects” from unobserved individual effects is common in the statistical survey literature. Clustering commonly arises in national field surveys from the fact that physically proximate households are often sampled to save time and money, but it can also arise from more homely sampling procedures. For example, Williams [2000; p.645] notes that it could arise from dental studies that “collect data on each tooth surface for each of several teeth from a set of patients” or “repeated measurements or recurrent events observed on the same person.” The procedures for allowing for clustering allow heteroskedasticity between and within clusters, as well as autocorrelation within clusters. They are closely related to the “generalized estimating equations” approach to panel estimation in epidemiology (see Liang and Zeger [1986]), and generalize the “robust standard errors” approach popular in econometrics (see Rogers [1993]). Wooldridge [2003] reviews some issues in the use of clustering for panel effects, in particular noting that significant inferential problems may arise with small numbers of panels.

²⁴ The data consists of observations from 253 subjects. Our data is a panel resulting in 7,928 risk aversion choices and 15,180 discount rate choices. In this specification we do not allow for any heterogeneity across subjects, so theory predicts a discontinuous jump in observed binary choices in the experiment. Since we pool over subjects with different preferences, we do not see this jump in the data; an appendix illustrates this pattern, and provides more details. We do account for observed and unobserved heterogeneity in the

the estimates allowing for risk aversion, and panel B presents the effects of constraining the model to assume risk neutrality. From panel A we obtain estimates of the CRRA parameter of 0.74. This estimate is slightly higher than 0.67 reported in HLRS (p.148). However, we use different statistical methods and the inclusion of background consumption in the utility function increases estimates of the CRRA parameter. The risk-neutral estimate of the discount rate is close to the estimate of 24.2% reported in HLRS (p. 151) using the same data but different statistical methods, and the 28.1% reported in HLW using a prior series of comparable field experiments in Denmark. Critically, allowing for risk aversion in panel A we obtain a point estimate of the discount rate of 10.1%, which is significantly lower than the estimate in panel B of 25.2%.

This result does more than simply verify that discount rates and risk aversion are mathematical substitutes in the sense that either of them has the effect of lowering the influence from future payoffs on present utility. It tells us that, for risk aversion coefficients that are reasonable from the standpoint of explaining choices in the lottery choice task, the estimated discount rate takes on a value that is more in line with what one would expect from market interest rates.²⁵ To evaluate the statistical significance of adjusting for a concave utility function we test the hypothesis that the estimated discount rate assuming risk aversion is the same as the discount rate estimated assuming risk neutrality. We easily reject this null hypothesis. Thus, *allowing for risk aversion makes a significant difference to the elicited discount rates.*²⁶

estimates presented in §3, and the same qualitative conclusions about discount rates are obtained.

²⁵ The discount rate is not the same thing as the intertemporal elasticity of substitution, which is the inverse of the risk aversion coefficient under the standard, temporally-separable model (e.g., Chavas [2004; pp. 145-146]). Using parlance from the general equilibrium calibration literature, the discount rate is the slope of the trade-off between temporally dated utility bundles, and not the elasticity of substitution between those bundles at some point. That elasticity is not constrained by the slope, *per se*.

²⁶ We do not vary the stakes in the earlier option across the discount rate tasks, and do not dramatically or proportionally vary them across the risk aversion tasks. If our subjects had actually faced tasks with dramatically lower or higher stakes, we would probably not have the same estimates that we report here. This is not to say that the estimates are fragile, just that if we varied the stakes over a significantly wider domain we would need to employ more flexible functional forms that would allow the risk parameter r to vary with the stake. Consider, for example, the existing evidence for the effects on estimated risk aversion as one moves from large stake tasks to small stake tasks. Holt and Laury [2002; p.1653] use an Expo-Power utility function and the estimate of relative risk aversion, r in our notation, varies from 0.27 at a prize level close to 0 up to 0.63 at a prize level of \$50, and then up to 0.87 at a prize level of \$100. Their estimates rest on monotonic and proportional changes in scale of rewards in experimental treatments, starting from a

The estimates of the two error terms in Table 3 exhibit an interesting pattern. Recall that estimates of zero indicate that no noise is present in the decision process, so we first observe that there is evidence of some noise since the p -values for each of ν and μ are statistically significant. However, there is a larger estimate of noise for the risk aversion tasks than the discount rate tasks in panel A, consistent with our prior that the risk aversion tasks were cognitively harder.

We can parametrically vary λ and evaluate the estimates of CRRA coefficients and discount rates. The left panel in Figure 1 shows that the CRRA coefficient is 0.74 when $\lambda=1$, and it increases to 0.75 when $\lambda=7$. Income from the risk aversion task is spent in one period and we therefore observe marginal changes in the CRRA coefficient when λ is varied. There is more variation in the estimated discount rate and the right panel shows that $\delta=10.1\%$ when $\lambda=1$, and that it increases to 14.5% when $\lambda=7$. The discount rate will continue to rise as λ increases, and it tends to the risk neutral level of 25.2% as λ goes to infinity. In fact, the log-likelihood of the model is maximized when $\lambda=1$, and steadily worsens as λ is incremented. The values of the likelihood are very similar for small λ , ranging from -12,459.12 at $\lambda=1$ to -12,459.37, -12,460.41, -12,461.81 and -12,463.38 at $\lambda=2, 3, 4$ and 5, respectively. By the time $\lambda=10$, for example, the likelihood has dropped to -12,471 and δ has risen to 16%. So there is some sensitivity of the estimates of δ to increments in the assumed λ , particularly for λ greater than a week, but the empirical evidence points to lower values of λ being appropriate. Estimated discount rates are still significantly lower than their risk neutral counterpart for λ up to two weeks ($\delta=17\%$).

Finally, we vary ω and evaluate the estimates of CRRA coefficients and discount rates when $\lambda=1$. The CRRA coefficient is 0.67 when $\omega=50$, and increases to 0.82 when $\omega=200$. There is very

baseline in which prizes are only \$0.10, \$1.60, \$2.00 and \$3.60, and then considering significant scale increases in prizes by factors of 20, 50 and 90. They implicitly assume a short-run self who is completely unconstrained, and do not integrate the money received with any background consumption. Nevertheless, these results indicate that changes in relative risk aversion can be quantitatively large when the changes in scale are by a common factor of 20 or more. Using the same data as we use, Harrison, Lau and Rutström [2007] found no statistically significant changes in estimated relative risk aversion over the domain of stakes in our tasks. This validates the use of the CRRA function for the stakes used here. However, that does not mean that we should assume CRRA for domains that we did not evaluate in our experiments. Thus we would expect some of the estimates in Table 3 to be different if the task given to subjects was significantly different.

little variation in the estimated discount rate as we vary ω : for example, we estimate $\delta=10.2\%$ when $\omega=50$, and $\delta=9.8\%$ when $\omega=200$. Hence, as we increase ω there is a slight increase in the difference between the estimated discount rate when we allow for concavity of the utility function and the estimated discount rate when we assume risk neutrality.

3. Alternative Specifications and Functional Forms

Although the basic insight that one should elicit risk and time preferences jointly seems simple enough, it is clear that identification of specific estimates does rely on assuming certain functional forms and statistical specifications. The specifications considered in section 2 and Table 3 are canonical, and important in their own right given their place in the literature. Can we say that our main conclusion is robust to alternative specifications and functional forms? One concern is with the effect of allowing for heterogeneity, in the sense that we allow estimable parameters to be linear functions of observable individual characteristics (section A) or reflect unobservable differences in preferences (section B). Another concern is with alternative discounting functions, such as those assumed in hyperbolic discounting models (section C). Finally, we consider statistical specifications that allow us to consider more than one latent data generating process and directly estimate the weight of the evidence favoring the different processes (section D).

A. Heterogeneity of Preferences and Observable Characteristics

It is an easy matter to allow each parameter in (11) to be a linear function of observable characteristics of individuals and/or treatment effects. For example, we could allow for the risk aversion parameter r to depend on the sex of the subject, so that we would estimate

$$\hat{r} = \hat{r}_0 + (\hat{r}_{\text{FEMALE}} \times \text{FEMALE}) \quad (12)$$

where \hat{r}_0 is the estimate of the constant and \hat{r}_{FEMALE} shows the difference in risk for females.

Allowing for demographic effects for r and δ makes no difference to our conclusion that allowing for risk aversion gives significantly lower estimates of discount rates. We include binary

indicators for sex, aged less than 30, aged between 40 and 50, aged over 50, living alone, having children, owning one's own home or apartment, being retired, being a student, having some post-secondary education, having substantial higher education, have a lower income level in 2002 (below 300,000 DKK), having a higher income level in 2002 (500,000 DKK or more), living in the greater Copenhagen area, and living in a larger city of 20,000 inhabitants or more. We also include a variable that measures the number of people in the household. Each of the core parameters r and δ is specified as a linear function of these characteristics, and the model is estimated using maximum likelihood. We also estimate the model assuming that everyone is risk-neutral and allowing for demographic heterogeneity in the estimate of δ .

The top panel A of Figure 2 displays kernel density estimates of the predicted discount rates from each estimation when $\omega=118$, $\eta=1$ and $\lambda=1$. When we assume that subjects are risk-neutral the mean discount rate is virtually identical to that estimated without controls for demographics. We test whether these demographic effects on discount rates are the same across the risk neutral and risk averse models, and confirm that they are not.²⁷ Not only is the mean of the predicted discount rate distribution for the risk averse case significantly below the risk neutral case, the variance is also much smaller. This lends further support for the value of estimating risk attitudes and discount rates simultaneously.

B. Heterogeneity of Preferences and Unobservable Characteristics

We can also extend the analysis to allow for heterogeneity of responses that are not correlated with observable individual characteristics. One way to do this is to view the core parameters r and δ as random, quite apart from any sampling error in estimation. Thus we presume that

$$r \sim N(r_m, r_{sd}^2) \tag{13}$$

²⁷ Further discussion of demographic effects on risk attitudes and discount rates in Denmark is provided by HLW, Harrison, Lau and Rutström [2007] and HLRS.

$$\delta \sim N(\delta_m, \delta_{sd}^2) \quad (14)$$

where r_m and δ_m are the mean values of the population parameters, r_{sd} and δ_{sd} are the standard deviations of the parameters over the population, and r and δ are assumed to be distributed according to a bivariate Normal distribution. Our previous estimates can be viewed as estimating r_m and δ_m under the assumption that $r_{sd}=\delta_{sd}=0$. We considered observable individual characteristics above, which of course goes a long way to recognize the population heterogeneity in these preference parameters. However, even after correcting for observable characteristics, there may be some unobservable heterogeneity left.

Estimation of random coefficients of this kind can be undertaken using maximum simulated likelihood (MSL) methods.²⁸ The essential idea of MSL methods is to simulate the likelihoods for random draws from the proposed distribution of r and δ and then average these simulated likelihoods. So each likelihood evaluation involves H extra evaluations, where H is the number of random draws from these distributions. Advances in computational power, and clever ways of drawing random sequences to ensure good coverage of the intended density with minimal H , make it feasible to undertake MSL for small-dimensional problems such as ours.

We find similar estimates of the heterogeneity of risk attitudes and discount rates, again setting $\omega=118$ DKK, $\eta=1$ and $\lambda=1$, and assuming zero-mean random effects from heterogeneity on r and δ . The estimated standard deviation r_{sd} is 0.056 and the estimated standard deviation δ_{sd} is 2.4%. If one imposes risk neutrality, the estimated standard deviation δ_{sd} is 17.1%. Panel B of Figure 2 displays the implied distributions of estimated discount rates. Thus, allowing for random coefficients has no major impact on our estimates of risk aversion and discount rates.

The bottom panel B of Figure 2 illustrates the relationship between the heterogeneity in discount rates when modeled using random coefficients and the alternative assumptions about risk

²⁸ Standard references include Gouriéroux and Monfort [1996], Train [2003] and Cameron and Trivedi [2005]. The basic numerical methods we employ were developed by Cappellari and Jenkins [2006] for *Stata*. We also implemented a truncated Normal distribution for δ to ensure that $\delta \geq 0$ (see Train [2003; §9.2.4]), and these estimates were virtually identical to those obtained by not truncating.

attitudes, similar to the manner in which panel A of Figure 2 shows this relationship based on observable demographics. The variances in both of the distributions in panel B are slightly greater than those in panel A, which should not surprise as the predictions in B account for sources of variation in behavior that are not captured by the observable characteristics in panel A. The increase in variability is particularly pronounced for the special case in which one imposes risk neutrality (the dashed lines), but is relatively slight for the general case in which one allows risk aversion (the solid lines). The main qualitative result on the importance of correcting for the concavity of the utility function is the same, irrespective of the two ways of correcting for heterogeneity.

C. Hyperbolic Discounting

The earliest hyperbolic specifications assumed that individuals had discount rates that declined with the horizon they faced, in contrast to later quasi-hyperbolic specifications that posit an initial decline and then constant (per period) discount rates. Our use of a front end delay on receipt of the earlier option implies that we cannot test the quasi-hyperbolic specification against the standard exponential specification, unless one assumes that the “passion for the present” lasted longer than our front end delay. We therefore focus on the earlier hyperbolic specifications. The most common functional form of the older literature is due to Herrnstein [1981], Ainslie [1992] and Mazur [1987], and would replace (7') and (8') with

$$PV_A = \sum_{i=\{t, \dots, t+\lambda-1\}} [(1/(1+\gamma \cdot i)) (\omega + M_A/\lambda)^{(1-\gamma)}] + \sum_{i=\{t+\tau, \dots, t+\tau+\lambda-1\}} [(1/(1+\gamma \cdot i)) \omega^{(1-\gamma)}] \quad (7'')$$

$$PV_B = \sum_{i=\{t, \dots, t+\lambda-1\}} [(1/(1+\gamma \cdot i)) \omega^{(1-\gamma)}] + \sum_{i=\{t+\tau, \dots, t+\tau+\lambda-1\}} [(1/(1+\gamma \cdot i)) (\omega + M_B/\lambda)^{(1-\gamma)}] \quad (8'')$$

for $\gamma > 0$, and with discounted utility to the time of decision, $t=0$.

Maximum likelihood estimates using (7'') and (8'') can be obtained using the same methods used for the exponential specification, and the results assuming $\lambda=1$ are reported in Table 4. We estimate γ to be 0.10 when one allows for risk aversion, and 0.27 when risk neutrality is imposed.

Figure 3 shows the implied hyperbolic estimates of annual discount rates against the time horizon in years. Clearly the hyperbolic model shows evidence of declining discount rates with

horizon when risk neutrality is assumed, but the quantitative magnitude of the decline is much smaller when one allows for concave utility functions. Furthermore, the level of discount rates is significantly lower in the latter case, which is consistent with our previous inferences based on exponential discounting.

We can also evaluate the effect of using a more general hyperbolic specification proposed by Prelec [2004; p.526]. This specification replaces (7') and (8') with

$$PV_A = \sum_{i=\{t, \dots, t+\lambda-1\}} [\exp\{-\beta \cdot i^\alpha\} (\omega + M_A/\lambda)^{(1-\tau)}] + \sum_{i=\{t+\tau, \dots, t+\tau+\lambda-1\}} [\exp\{-\beta \cdot i^\alpha\} \omega^{(1-\tau)}] \quad (7''')$$

$$PV_B = \sum_{i=\{t, \dots, t+\lambda-1\}} [\exp\{-\beta \cdot i^\alpha\} \omega^{(1-\tau)}] + \sum_{i=\{t+\tau, \dots, t+\tau+\lambda-1\}} [\exp\{-\beta \cdot i^\alpha\} (\omega + M_B/\lambda)^{(1-\tau)}] \quad (8''')$$

The exponential discounting model emerges as a limiting case as α tends to 1. One can think of the parameter α as characterizing the “decreasing impatience” of the decision-maker, a smoother and inverse counterpart of the notion of a “passion for the present” in quasi-hyperbolic specifications. As α takes values below 1, the discount function takes on the familiar shape of earlier hyperbolic specifications. One can also think of the parameter β as characterizing time preferences in the usual sense (Prelec [2004; p.524]). The instantaneous discount rate implied by this discount function is $\alpha\beta t^{\alpha-1}$, which collapses to β as $\alpha \rightarrow 1$. If we use this specification (8''') instead of the exponential discounting specification (8') or the hyperbolic specification (8'') we obtain essentially the same results: elicited discount rates are significantly lower when one allows for concave utility functions.

D. Mixture Specifications

Finally, we consider the sensitivity of our conclusions to a statistical specification that allows each observation to potentially be generated by more than one latent data-generating process. Our motivation is to better identify the effect of jointly eliciting risk and time preferences on the weight of evidence for exponential and hyperbolic discounting. Consider Figure 3 again, which assumes that all of the data was generated by one hyperbolic data-generating process. The significant decline in discount rates under risk neutrality is muted when one accounts for risk aversion. How can we use these data to inform us about the relative importance of the exponential and hyperbolic specification

after we allow for concave utility functions?

Finite mixture models provide an ideal statistical framework to address this question.²⁹

Consider the mixture of exponential discounting models and hyperbolic models, defined by (8') and (8''). We assume that EUT characterizes behavior in all other respects. The mixture likelihood function is then

$$\ln L(r, \delta, \gamma, \mu, \nu, \pi; y, \omega, \lambda, \mathbf{X}) = \ln L^{\text{RA}} + [\pi \times \ln L^{\text{DR-E}}] + [(1-\pi) \times \ln L^{\text{DR-H}}] \quad (15)$$

where π is a parameter, to be estimated and constrained such that $0 \leq \pi \leq 1$, giving the probability that a given observation³⁰ is generated by the exponential discounting model. In (15) the likelihood contributions $L^{\text{DR-E}}$ and $L^{\text{DR-H}}$ refer to the exponential and hyperbolic specifications given by (8') and (8''), respectively.

Table 5 provides maximum likelihood estimates of this model, where estimates allowing for a concave utility function are presented in Panel A and risk neutral estimates are presented in Panel B. We again assume that $\lambda=1$ in both the exponential and the hyperbolic specifications. Figure 4 displays the predicted discounting functions for each model. From Panel A in Table 5 we see that, when we allow for a concave utility function, 72% of the observations can be characterized as being generated by the exponential specification with a discount rate of 6.8%, and the remaining 28% of the observations are generated by the hyperbolic specification with $\gamma=0.33$. This value of γ implies the discounting function shown in the top, right panel of Figure 4: discount rates around 50% p.a for a 3-month horizon, and around 20% p.a. for a 1-year horizon.

This mixture specification suggests that assuming that behavior is completely exponential *or*

²⁹ Mixture models have an astonishing pedigree in statistics: Pearson [1894] examined data on the ratio of forehead to body length of 1000 crabs to illustrate “the dissection of abnormal frequency curves into normal curves,...”. In modern parlance he was allowing the observed data to be generated by two distinct Gaussian processes, and estimated the two means and two standard deviations. Modern surveys of the evolution of mixture models are provided by Everitt [1996] and McLachlan and Peel [2000]. Harrison and Rutström [2005] review the literature on mixture models in experimental economics, and discuss the interpretation of alternative mixture specifications.

³⁰ One could alternatively define a grand likelihood in which observations or subjects are completely classified as following one model or the other on the basis of the latent probability π . El-Gamal and Grether [1995] illustrate this approach in the context of identifying behavioral strategies in Bayesian updating experiments.

hyperbolic is in error. Both specifications have some support, even if one has greater support than the other. It also suggests that the earlier finding that the discount rate was 10.1% in section 2 masked two very different latent processes at work. It picked up the 6.8% that characterized 72% of the sample in an exponential manner, but it also factored in the 28% of the sample that had significantly higher discount rates for shorter horizons. In other words, it was a weighted average of the two top panels of Figure 4, “forced” into an exponential functional form. Of course, the same is true of the hyperbolic specification assumed in the right panel of Figure 3.

Turning to the risk neutral estimates in Panel B we observe that 93% of the observations are generated by the exponential specification with a discount rate of 23.3%, and 7% of the observations can be explained by the hyperbolic specification with $\gamma=1.55$. This value of γ implies the relatively steep discounting function shown in the bottom, right panel of Figure 4. Discount rates are around 150% p.a. for a 3-month horizon, and still over 50% p.a. for a 6-month horizon. Although the hyperbolic specification garners a much smaller fraction of the sample when one assumes risk neutrality, it is associated with very high discount rates for the shorter horizons. Hence, the hyperbolic model captures choices that are associated with extreme levels of discount rates while the majority of the choices indicate constant and moderate levels of discount rates.

4. Related Literature

There are several studies that note the connection between concave utility functions and individual discount rates, but we are aware of only two studies that address the joint elicitation of risk and time preferences directly using monetary incentives.³¹

A. Studies Using Hypothetical Tasks

Chapman [1996] draws the correct formal link between estimation of individual discount

³¹ There are some studies that undertake joint statistical estimation of discount rates and risk attitudes from field data or survey data. For example, see Abdulkadri and Langemeier [2000], Issler and Piquera [2000] and van Praag and Booij [2003].

rates and concavity of the utility function, but does not elicit risk attitudes. Instead she uses hypothetical questions to elicit individual discount rates over money and health, and then estimates individual discount rates based on various assumptions about the risk attitudes of the subjects.

Kirby and Santiesteban [2003] argue correctly that the hyperbolic behavior found in many studies may be confounded by the concavity of the utility function and transactions costs with respect to the delayed payment option. However, they do not control for risk attitudes in their evaluation of different discounting functions.

B. Studies Using Monetary Incentives

Anderhub, Güth, Gneezy and Sonsino [2001] (AGGS) use the Becker-DeGroot-Marschak (BDM) procedure to elicit certainty equivalents for lotteries with varied payoff dates.³² They used undergraduate economic students in Israel as subjects. Each subject provided either a buying or a selling price for each of three lotteries that paid out the day of the experiment, two weeks from the day of the experiment, and four weeks from the day of the experiment. The lotteries differ only with respect to the timing of payments. One decision was chosen at random to be played out. AGGS find no statistical difference between certainty equivalents across different time horizons. They find a significant *positive* relationship between the degree of risk aversion and the discount rates implied by the timing of payments.³³ The differences between the elicitation tasks in our design and that of AGGS reflect a tradeoff between compactness of experimental procedures and transparency of the task required of subjects. While our elicitation mechanism is logically equivalent to the BDM, we believe the binary decisions in the MPL are less of a cognitive burden for subjects. Moreover, the AGGS design elicits a single value from subjects that reflects both risk and time preferences, while we examine these preferences separately.

³² Ahlbrecht and Weber [1997] use a similar design, albeit with hypothetical rewards, and find no significant relationship between risk aversion and individual discount rates in the gain domain.

³³ AGGS are able to effect delayed payments by distributing post-dated checks the day of the experiment, thereby reducing any differences between immediate and delayed payments due to subject expectations regarding transactions costs of future payments.

The main difference between our results and those of AGGS is that we use the information on risk attitudes to infer the discount rate defined over utility, rather than the discount rate over money. Using the specification of §3.A in which we allow for observable demographics characteristics, we can predict risk attitudes and discount rates for each of our subjects. The rank correlation of those predicted values is 0.10 with a p -value of 0.021, so we also find evidence of a positive correlation between risk aversion and impatience.³⁴ The rank correlation is 0.32, with a p -value less than 0.001, if one instead estimates the discount rates under the assumption that the subject is risk neutral. So correcting for the fact that discount rates are defined over utility streams not just reduces the estimated discount rates, but also reduces the correlation between risk aversion and impatience.

Eckel, Johnson and Montmarquette [2005] (EJM) conduct a field study of time and risk preferences. Their subjects are recruited from low income neighborhoods in Montreal. Subjects in these experiments are given 64 “compensated” questions, one of which is chosen at random for payment. Time preferences are elicited by presenting subjects with choices between payoffs that occur at different times. Time horizons for the later payments ranged from 2 days to 28 days, and most early payments had a front end delay of one day, one week, or two weeks. The value for most questions started at approximately \$72 CAD, with a few questions presenting values around \$26 CAD. The distribution of annual discount rates implied by the questions was lumpy, with values of 10%, 50%, 200% and 380%. In fact, EJM [2005; p.258] report short-term discount rates *averaging* 289% per annum, consistent with the earlier literature on discount rate elicitation.³⁵ Given the variance of responses, EJM construct a measure of time preferences that rely on four questions

³⁴ This correlation across subjects is unrelated to the mathematical fact that the observed choices of a *given subject* would imply lower discount rates if we assume that the subject was risk averse instead of assuming that the subject was risk neutral. Of course, allowing for this fact is critical to estimating the correct correlations across subjects.

³⁵ Rates as high as this are actually quite common in the extensive psychology literature where procedures are quite different from ours (e.g., Kirby and Maraković [1996]). The use of hypothetical scenarios, scrambling of the ordering of choices, and the absence of information on interest rates are common. Following CW, who also review earlier economics experiments that do not use hypothetical scenarios, it is now common to present subjects with an ordered series of choices to reduce simple confusion, and to present the interest rate information.

using the 14 days time horizon only. Risk preferences are elicited in a similar fashion by presenting subjects with choices between lotteries in random order, where most choices involved a “less risky” lottery that paid a single amount with certainty.³⁶ The expected value of the lotteries ranged from \$40 CAD to \$120 CAD. EJM find that subjects who choose the less risky lotteries have significantly higher individual discount rates, which is also consistent with our finding, but they do not estimate the relationship between risk and time preferences, or consider alternative functional forms. Again, their finding refers to discount rates defined over money, which overstates the correct correlation between risk attitudes and discount rates defined over utility.

Tanaka, Camerer and Nguyen [2007] propose methods to elicit risk preferences and time preferences from individuals. Harrison and Rutström [2007; §1.6] discuss the logic and maintained assumptions of their approach. However, their elicitation procedure for time preferences is conceptually completely separate from their elicitation procedure for risk attitudes, and responses to the risk attitude elicitation tasks are not used to condition inferences about time preferences. In effect, their procedure assumes risk neutrality when inferring discount rates, which will lead to overestimates of discount rates between utility flows since their subjects exhibit risk aversion. On the other hand, one might be able to use the raw responses to their tasks to undertake joint inference about risk and time preferences, using the maximum likelihood methods proposed here.

5. Conclusions

We find that credible estimates of discount rates rely on the joint estimation of risk and time preferences. When one assumes that subjects are risk neutral when in fact they are risk averse, the estimated discount rates are significantly biased upwards. Most of the literature has been aware of this bias, but until now has not been able to formalize the estimation of time preferences to the point of decomposing it. We also show that this effect is independent of the adoption of alternative

³⁶ Each of the first ten questions presented two lotteries with the same expected value, while the expected value of the less risky lottery was lower than the expected value of the more risky lottery in the last four questions. The predicted value of CRRA is 0.78, which is similar to other experimental evidence.

specifications, including hyperbolic discounting, observable and unobservable heterogeneity and the possibility of several latent data generating processes. This robustness across specifications is important given the role of parametric theory in allowing the identification of latent time preferences from observed choices.

Our results have direct implications for future efforts to elicit time preferences. The obvious one is to jointly elicit risk and time preferences, or at least to elicit risk preferences from a sample drawn from the same population, so that inferences about time preferences can be conditioned appropriately. There are also broader implications for the testing of theories of choice over time. Many “discounting anomalies” have been suggested in the literature, and it is unclear *a priori* how the proper accounting for concave utility functions affects these anomalies.

References

- Abdulkadri, A.O., and Langemeier, M.R., "Using Farm Consumption Data to Estimate the Intertemporal Elasticity of Substitution and Relative Risk Aversion Coefficients," *Agricultural Finance Review*, 60, 2000, 61-70.
- Ahlbrecht, Martin, and Weber, Martin, "An Empirical Study on Intertemporal Decision Making Under Risk," *Management Science*, 43(6), June 1997, 813-826.
- Ainslie, George, *Picoeconomics* (New York: Cambridge University Press, 1992).
- 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; Harrison, Glenn W.; Lau, Morten I., and Rutström, E. Elisabet, "Elicitation Using Multiple Price Lists," *Experimental Economics*, 9(4), December 2006, 383-405.
- Andersen, Steffen; Harrison, Glenn W., Lau, Morten I., and Rutström, E. Elisabet, "Lost in State Space: Are Preferences Stable?" *International Economic Review*, 49, 2008 forthcoming.
- Ballinger, T. Parker, and Wilcox, Nathaniel T., "Decisions, Error and Heterogeneity," *Economic Journal*, 107, July 1997, 1090-1105.
- Barrett, Lisa Feldman; Tugade, Michele M., and Engle, Randall W., "Individual Differences in Working Memory Capacity and Dual-Process Theories of the Mind," *Psychological Bulletin*, 130(4), 2004, 553-573.
- Benhabib, Jess, and Bisin, Alberto, "Modeling internal commitment mechanisms and self-control: A neuroeconomics approach to consumption-saving decisions," *Games and Economic Behavior*, 52, 2005, 460-492.
- Cameron, A. Colin, and Trivedi, Pravin K., *Microeconometrics: Methods and Applications* (New York: Cambridge University Press, 2005).
- Cappellari, Lorenzo, and Jenkins, Stephen P., "Calculation of Multivariate Normal Probabilities by Simulation, with Applications to Maximum Simulated Likelihood Estimation," *Stata Journal*, 6(2), 2006, 156-189.
- Chapman, Gretchen, "Temporal Discounting and Utility for Health and Money," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(3), 1996, 771-791.
- Chavas, Jean-Paul, *Risk Analysis in Theory and Practice* (San Diego, CA: Elsevier Academic Press, 2004).
- Cohen, Jonathan D., "The Vulcanization of the Human Brain: A Neural Perspective on Interactions Between Cognition and Emotion," *Journal of Economic Perspectives*, 19(4), Fall 2005, 3-24.
- Coller, Maribeth, and Williams, Melonie B., "Eliciting Individual Discount Rates," *Experimental Economics*, 2, 1999, 107-127.
- 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).
- El-Gamal, Mahmoud A., and Grether, David M., "Are People Bayesian? Uncovering Behavioral Strategies," *Journal of the American Statistical Association*, 90, 432, December 1995, 1137-1145.
- Everitt, Brian S., "An Introduction to Finite Mixture Distributions," *Statistical Methods in Medical Research*, 5, 1996, 107-127.
- Frederick, Shane; Loewenstein, George, and O'Donoghue, Ted, "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature*, XL, 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.
- Fudenberg, Drew, and Levine, David K., "Self Control, Risk Aversion and the Allais Paradox," *Unpublished Manuscript*, Department of Economics, Washington University at St. Louis, 2007.
- Gollier, Christian, *The Economics of Risk and Time* (Cambridge, MA: MIT Press, 2001).
- Gouriéroux, Christian, and Monfort, Alain, *Statistics and Econometric Models* (Cambridge, UK: Cambridge University Press, 1996).
- Hamermesh, Daniel S., and Pfann, Gerard A. (eds.), *The Economics of Time Use* (Amsterdam: Elsevier, Contributions to Economic Analysis #271, 2005).
- Harless, David W., and Camerer, Colin F., "The Predictive Utility of Generalized Utility Theories," *Econometrica*, 62(6), November 1994, 1251-1289.
- Harrison, Glenn W., Harstad, Ronald M., and Rutström, E. Elisabet, "Experimental Methods and Elicitation of Values," *Experimental Economics*, 7(2), June 2004, 123-140.
- 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 I., and Rutström, E. 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, "Expected Utility Theory *And* Prospect Theory:

- One Wedding and A Decent Funeral,” *Working Paper 05-18*, Department of Economics, College of Business Administration, University of Central Florida, 2005.
- 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* (Greenwich, CT: JAI Press, Research in Experimental Economics, Volume 12, 2007 forthcoming).
- Herrnstein, Richard J., “Self-Control as Response Strength,” in C. M. Bradshaw, E. Szabadi and C.F. Lowe (eds.), *Quantification of Steady-State Operant Behavior* (Amsterdam: North-Holland, 1981).
- Hey, John D., “Experimental Investigations of Errors in Decision Making Under Risk,” *European Economic Review*, 39, 1995, 633-640.
- Hey, John D., “Experimental Economics and the Theory of Decision Making Under Uncertainty,” *Geneva Papers on Risk and Insurance Theory*, 27(1), June 2002, 5-21.
- Hey, John D., and Orme, Chris, “Investigating Generalizations of Expected Utility Theory Using Experimental Data,” *Econometrica*, 62(6), November 1994, 1291-1326.
- 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.
- Issler, J. V., and Piqueira, N. S., “Estimating Relative Risk Aversion, the Discount Rate, and the Intertemporal Elasticity of Substitution in Consumption for Brazil Using Three Types of Utility Functions,” *Brazilian Review of Econometrics*, 20, 2000, 201-239.
- Kahneman, Daniel; Knetsch, Jack L., and Thaler, Richard H., “Experimental Tests of the Endowment Effect and the Coase Theorem,” *Journal of Political Economy*, 98, December 1990, 1325-1348.
- Keller, L. Robin, and Strazzera, Elisabetta, “Examining Predictive Accuracy Among Discounting Models,” *Journal of Risk and Uncertainty*, 24(2), 2002, 143-160.
- 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., 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.
- Kreps, David M., and Porteus, Evan L., “Temporal Resolution of Uncertainty and Dynamic Choice Theory,” *Econometrica*, 46(1), January 1978, 185-200.
- Liang, K-Y, and Zeger, S. L., “Longitudinal Data Analysis Using Generalized Linear Models,” *Biometrika*, 73, 1986, 13-22.
- Loomes, Graham; Moffatt, Peter G., and Sugden, Robert, “A Microeconomic Test of Alternative Stochastic Theories of Risky Choice,” *Journal of Risk and Uncertainty*, 24(2), 2002, 103-130.

- Loomes, Graham, and Sugden, Robert, "Incorporating a Stochastic Element Into Decision Theories," *European Economic Review*, 39, 1995, 641-648.
- 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).
- McLachlan, Geoffrey, and Peel, David, *Finite Mixture Models* (New York: Wiley, 2000).
- Mitchell, Robert C., and Carson, Richard T., *Using Surveys to Value Public Goods: The Contingent Valuation Method* (Baltimore: Johns Hopkins Press, 1989).
- Papke, Leslie E., and Wooldridge, Jeffrey M., "Econometric Methods for Fractional Response Variables with an Application to 401(K) Plan Participation Rates," *Journal of Applied Econometrics*, 11, 1996, 619-632.
- Pearson, Karl, "Contribution to the mathematical theory of evolution," *Philosophical Transactions A*, 185, 1894, 71-110.
- Prelec, Drazen, "Decreasing Impatience: A Criterion for Non-stationary Time Preference and "Hyperbolic" Discounting," *Scandinavian Journal of Economics*, 106(3), 2004, 511-532.
- Rabin, Matthew, "Risk Aversion and Expected Utility Theory: A Calibration Theorem," *Econometrica*, 68, 2000, 1281-1292.
- Rogers, W. H., "Regression Standard Errors in Clustered Samples," *Stata Technical Bulletin*, 13, 1993, 19-23.
- Shapiro, Jesse M., "Is There a Daily Discount Rate? Evidence from the Food Stamp Nutrition Cycle," *Journal of Public Economics*, 89, 2005, 303-325.
- Tanaka, Tomomi; Camerer, Colin F., and Nguyen, Quang, "Risk and Time Preferences: Experimental and Household Survey Data from Vietnam," *Working Paper*, California Institute of Technology, June 2007.
- Train, Kenneth E., *Discrete Choice Methods with Simulation* (New York: Cambridge University Press, 2003).
- van Praag, Bernard, M. S., and Booij, Adam S., "Risk Aversion and the Subjective Time Preference: A Joint Approach," *CEPR Working Paper 923*, Ifo Institute for Economic Studies, Munich, 2003.
- Wakker, Peter P., "Explaining the Characteristics of the Power (CRRA) Utility Family," *Unpublished Manuscript*, Econometric Institute, Erasmus University Rotterdam, 2006.
- Williams, Rick L., "A Note on Robust Variance Estimation for Cluster-Correlated Data," *Biometrics*, 56, June 2000, 645-646.
- Wooldridge, Jeffrey, "Cluster-Sample Methods in Applied Econometrics," *American Economic Review (Papers & Proceedings)*, 93(2), May 2003, 133-138.

Table 1: Typical Payoff Matrix in the Risk Aversion Experiments

Lottery A			Lottery B			EV ^A	EV ^B	Difference	Open CRRA Interval if Subject Switches to Lottery B and $\omega=0$
P	DKK	P	P	DKK	P				
0.1	2000	0.9	0.1	3850	0.9	100	475	1165	$-\infty, -1.71$
0.2	2000	0.8	0.2	3850	0.8	100	850	830	-1.71, -0.95
0.3	2000	0.7	0.3	3850	0.7	100	1225	495	-0.95, -0.49
0.4	2000	0.6	0.4	3850	0.6	100	1600	160	-0.49, -0.15
0.5	2000	0.5	0.5	3850	0.5	100	1975	-175	-0.15, 0.14
0.6	2000	0.4	0.6	3850	0.4	100	2350	-510	0.14, 0.41
0.7	2000	0.3	0.7	3850	0.3	100	2725	-845	0.41, 0.68
0.8	2000	0.2	0.8	3850	0.2	100	3100	-1180	0.68, 0.97
0.9	2000	0.1	0.9	3850	0.1	100	3475	-1515	0.97, 1.37
1	2000	0	1	3850	0	100	3850	-1850	1.37, ∞

Note: The last four columns in this table, showing the expected values of the lotteries and the implied CRRA intervals, were not shown to subjects.

Table 2: Payoff Table for 6 Month Time Horizon in the Discount Rate Experiments

Payoff Alternative	Payment Option A (pays amount below in 1 month)	Payment Option B (pays amount below in 7 months)	Annual Interest Rate (AR, in percent)	Annual Effective Interest Rate (AER, in percent)	Preferred Payment Option (Circle A or B)	
1	3,000 DKK	3,075 DKK	5	5.09	A	B
2	3,000 DKK	3,152 DKK	10	10.38	A	B
3	3,000 DKK	3,229 DKK	15	15.87	A	B
4	3,000 DKK	3,308 DKK	20	21.55	A	B
5	3,000 DKK	3,387 DKK	25	27.44	A	B
6	3,000 DKK	3,467 DKK	30	33.55	A	B
7	3,000 DKK	3,548 DKK	35	39.87	A	B
8	3,000 DKK	3,630 DKK	40	46.41	A	B
9	3,000 DKK	3,713 DKK	45	53.18	A	B
10	3,000 DKK	3,797 DKK	50	60.18	A	B

Table 3: Estimates of Risk and Time Preferences Assuming Exponential Discounting

Parameter	Estimate	Standard Error	Lower 95% Confidence Interval	Upper 95% Confidence Interval
<i>A. Allowing a Concave Utility Function (Risk Aversion)</i>				
r	0.741	0.048	0.648	0.835
δ	0.101	0.008	0.084	0.117
μ (for RA)	0.086	0.015	0.056	0.116
v (for DR)	0.023	0.005	0.012	0.034
<i>B. Assuming a Linear Utility Function (Risk Neutrality)</i>				
δ	0.252	0.012	0.228	0.276
v (for DR)	0.133	0.008	0.118	0.148

Figure 1: Estimated Discount Rates and Days of Consumption




Figure 2: Estimated Discount Rates
Allowing for Heterogeneity of Preferences
Assuming exponential discounting and EUT





Table 4: Estimates of Risk and Time Preferences Assuming Hyperbolic Discounting

Parameter	Estimate	Standard Error	Lower 95% Confidence Interval	Upper 95% Confidence Interval
<i>A. Allowing a Concave Utility Function (Risk Aversion)</i>				
r	0.750	0.048	0.656	0.844
γ	0.103	0.009	0.085	0.120
μ (for RA)	0.084	0.015	0.054	0.113
ν (for DR)	0.022	0.005	0.012	0.033
<i>B. Assuming a Linear Utility Function (Risk Neutrality)</i>				
γ	0.270	0.015	0.241	0.298
ν (for DR)	0.136	0.008	0.120	0.152

Figure 3: Assuming Hyperbolic Discounting





Table 5: Mixture Model Estimates of Risk and Time Preferences

Parameter	Estimate	Standard Error	Lower 95% Confidence Interval	Upper 95% Confidence Interval
<i>A. Allowing a Concave Utility Function (Risk Aversion)</i>				
r	0.774	0.048	0.680	0.869
δ	0.068	0.008	0.052	0.083
γ	0.333	0.040	0.256	0.411
μ (for RA)	0.077	0.015	0.046	0.107
v (for DR)	0.012	0.004	0.005	0.019
π	0.720	0.042	0.638	0.803
<i>B. Allowing a Linear Utility Function (Risk Neutrality)</i>				
δ	0.233	0.020	0.194	0.272
γ	1.550	0.276	1.009	2.091
v (for DR)	0.121	0.018	0.087	0.155
π	0.926	0.064	0.800	1.051

Figure 4: Mixture Specifications of Exponential and Hyperbolic Discounting Models

Estimated assuming linear utility or allowing concave utility



Appendix A: Actual and Implied Choices

Figure A1 shows implied choices in the discounting tasks, given the default point estimates of our model from Table 3. The vertical axis shows the fraction of responses in which option B was chosen, and the horizontal axis lists the option B values. Obviously these values get larger as the horizon increases. The solid line in each panel is the average of observed choices over all subjects in the sample. The dashed line in each panel is the implied switch point of the model evaluated at the point estimate.

The implied choices in Figure A1 are based on the point estimates from the homogeneous preferences model of Table 3. It assumes every agent has the same preferences, does not account for the standard error of the point estimate, and is therefore implies a “sharp” choice path. One could, in principle, generate a series of implied choices by bootstrapping these estimates around the point estimates, and then the implied choices would look “smoother.” This would capture the sampling error displayed in Table 3, but not the heterogeneity across subjects.

Heterogeneity in implied responses across subjects is captured by using the estimates generated by the model which controls for observable demographic characteristics of the individual subjects. This is the model estimated for Panel A of Figure 2. Doing this generates the implied choices in Figure A2, which are much smoother since they reflect differences in predictions for each subject (reflecting their demographic characteristics).

We could extend this exercise to use the estimates and implied choices from the model used to generate Panel B of Figure 2, allowing for unobserved heterogeneity, but the broad descriptive ability of the model is already apparent from Figure A2.

Figure A1: Actual and Implied Choices
Assuming Homogeneous Preferences
Fraction of Choices of B for each Option B Value (in DKK)

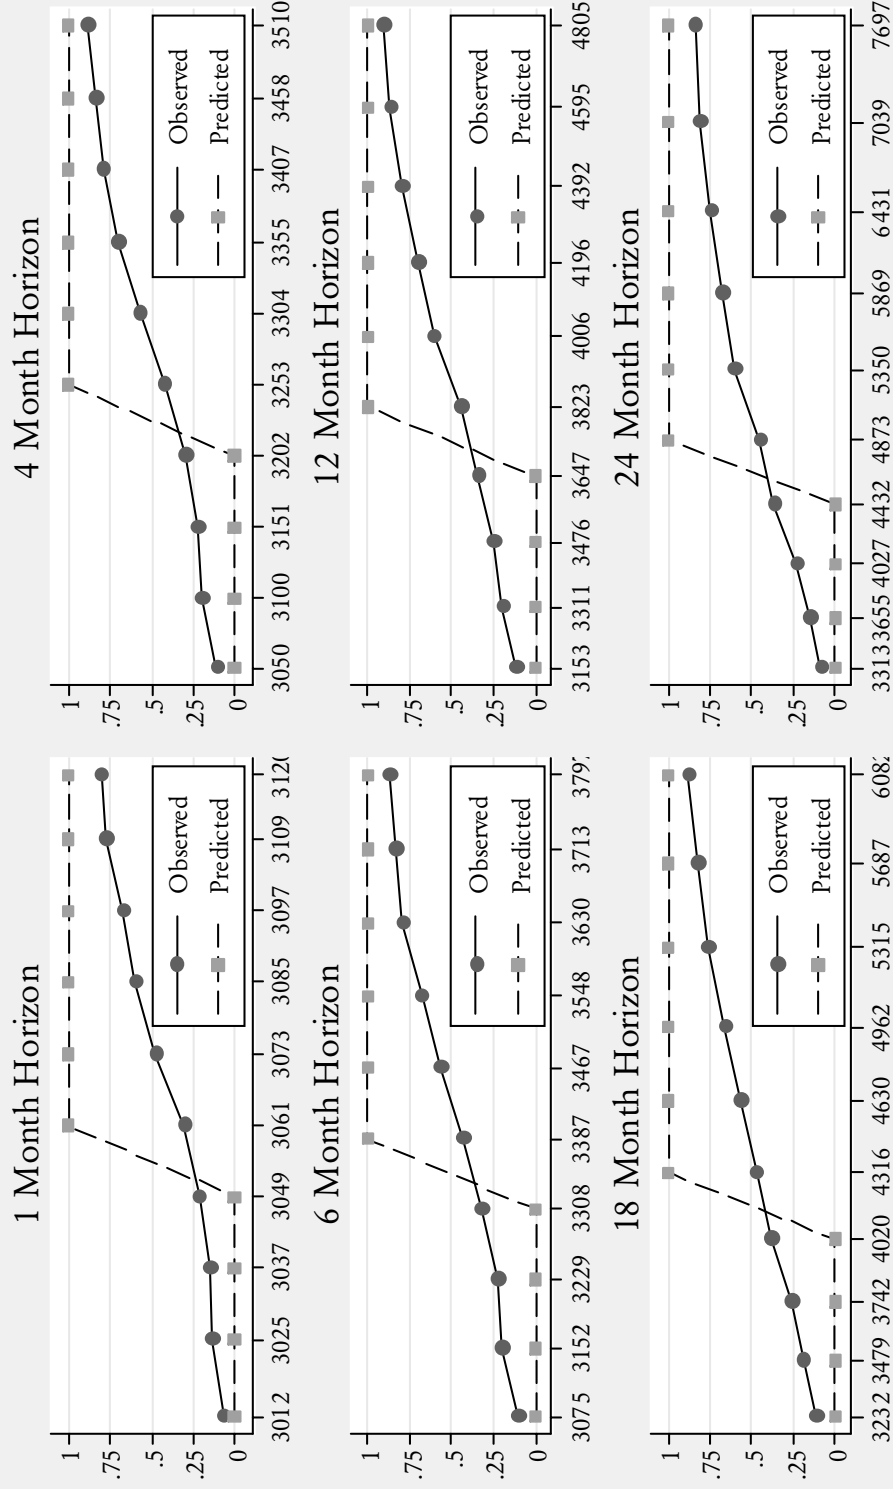


Figure A2: Actual and Implied Choices
Assuming Heterogeneous Preferences
Fraction of Choices of B for each Option B Value (in DKK)

