European Economic Review 71 (2014) 15–33

Contents lists available at ScienceDirect

European Economic Review

journal homepage: www.elsevier.com/locate/eer

Discounting behavior: A reconsideration ☆

Steffen Andersen [a], Glenn W. Harrison [b,*], Morten I. Lau [a,c], E. Elisabet Rutström [d]

[a] Department of Economics, Copenhagen Business School, Copenhagen, Denmark
[b] Department of Risk Management & Insurance and Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University, USA
[c] Durham University Business School, Durham University, UK
[d] Dean's Behavioral Economics Laboratory, Robinson College of Business, and Andrew Young School of Policy Studies, Georgia State University, USA

ARTICLE INFO

Article history:
Received 18 April 2013
Accepted 6 June 2014
Available online 16 July 2014

Keywords:
Discounting behavior
Field experiments
Behavioral microeconomics

ABSTRACT

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

1. Introduction

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

☆ We thank an Associate Editor and three referees for helpful, detailed comments, and the U.S. National Science Foundation for research support under grants NSF/HSD 0527675 and NSF/SES 0616746, the Danish Social Science Research Council for research support under project 275-08-0289, and the Carlsberg Foundation under grant 2008_01_0410. Additional appendices are available in *CEAR Working Paper 2011-03* available at http://cear.gsu.edu/papers.

* Corresponding author.
E-mail addresses: sa.eco@cbs.dk (S. Andersen), gharrison@gsu.edu (G.W. Harrison), mla.eco@cbs.dk (M.I. Lau), erutstrom@gsu.edu (E.E. Rutström).

http://dx.doi.org/10.1016/j.euroecorev.2014.06.009
0014-2921/

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

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

In Section 2 we review the alternative theoretical models that have been proposed, and settle on a list of major exemplars of the different types of models. We assume expected utility theory (EUT) for this initial purpose, and we consider alternatives as a robustness check later.[1] In Section 3 we use these theoretical structures to guide the design of a series of experiments that will allow us to identify the core parameters of the latent structural models. We also discuss our specific experiments, conducted throughout Denmark in 2009 using a representative sample of the adult Danish population. In Section 4 we review the econometric models used to estimate the core parameters of the models. We also explain how finite mixture models can be used to evaluate the heterogeneity of discounting behavior in the population. Section 5 contains basic results, and explores variations in some of the maintained assumptions of our basic results.

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

Given the contrary nature of our findings, in terms of the received empirical wisdom, Section 6 contains a systematic cataloguing of the samples, experimental procedures, and econometric procedures of the evidence for quasi-hyperbolic and non-constant discounting. We conclude that the evidence needed reconsideration.

One important robustness check is to see if the lack of support for the quasi-hyperbolic model is attributable to our population being the entire adult Danish population, rather than university students. Although it is apparent that the wider population is typically of greater interest, virtually all prior experimental evidence that we give credence to comes from convenience samples of university students. We find that there is indeed a difference in the elicited discount rates with (Danish) university students, but that they do not exhibit statistically significant evidence of declining discount rates. The size of the discount rates for shorter time horizons is greater than that of the general population, but much smaller than the received wisdom suggests.

2. Theory

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

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

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

for $t \geq 0$, and where the discount rate d is simply $d^E(t) = \delta$. Although these characterizations are abstract, we view the discount rate on an annualized basis throughout. The key feature of this model, of course, is that the discount rate is a constant over time. The percentage rate at which utility today and utility tomorrow is discounted is exactly the same as the rate at which utility in 7 days and utility in 8 days is discounted. The debate over climate change has reminded us all that,

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

with this specification, even small discount rates can lead to very low weight being placed on longer-term future consequences.

The discount factor for the quasi-hyperbolic (QH) specification is defined as

$$D^{\mathrm{QH}}(t) = 1 \quad \text{if } t = 0 \tag{2a}$$

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

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

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

$$D^{\mathrm{FC}}(t) = 1 \quad \text{if } t = 0 \tag{3a}$$

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

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

There have been whole families of "hyperbolic" specifications of the discounting function. The simplest specification assumes a discount factor given by $D^{\mathrm{H1}}(t) = 1/t$ with discount rates $d^{\mathrm{H1}}(t) = t^{(1/t)} - 1$. The H1 specification was proposed in this manner by Ainslie (1975; p. 472, Fig. 3). A slight generalization is

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

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

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

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

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

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

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

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

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

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

3. Experiments

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

3.1. Essential characteristics of the experiments

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

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

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

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

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

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

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

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

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

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

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

3.2. *The experimental design*

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

3.2.1. *Individual discount rates*

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

We ask subjects to evaluate choices over several time horizons. We consider time horizons between 2 weeks and 1 year. Each subject is presented with choices over four time horizons, and those horizons are drawn at random, without replacement, from a set of thirteen possible horizons (2 weeks, and 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11 and 12 months). This design will allow us to obtain a smooth characterization of the discount rate function across the sample for horizons up to one year. We also over-sampled the first three horizons, since this very short-term is clearly of great significance for the alternative specification. Hence each subject was twice as likely to get a horizon of 2 weeks, 1 month or 2 months as any of the later horizons.[9]

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

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

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

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

end delay treatment. We also varied the order in which the time horizon was presented to the subject: either in ascending order or descending order.

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

These four treatments, the front end delay, information on implied interest rates, the level of the principal, and the order of presentation of the horizon, gives a $2 \times 2 \times 2 \times 2$ design. The subjects were assigned at random to any one particular combination of treatments, and the sample is roughly split across the 16 treatments.

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

3.2.2. *Risk attitudes*

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

We assume that the utility function is stable over time and is perceived *ex ante* to be stable over time. Direct evidence for the former proposition is provided by Andersen et al. (2008b), who examine the temporal stability of risk attitudes in the Danish population. The second proposition is a more delicate matter. Even if utility functions are stable over time, they may not be subjectively perceived to be, and that is what matters for us to assume that it is the same utility function that appears in the definition of the discount factor. When there is no front end delay, this assumption is immediate for the smaller-sooner amount y, but not otherwise. Whether or not individuals suffer from a "projection bias" is a deep matter, demanding more research: see Ainslie (1992; p. 144–179, Section 6.3), Kirby and Guastello (2001) and Loewenstein et al. (2003).

We also assume that the same utility function that governs decisions over risky alternatives is the one that is used to evaluate time-discounted choices. This assumption has been criticized recently, and we take up those issues in Section 7.

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

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

3.3. *The experiments*

Between September 28 and October 22, 2009, we conducted experiments with 413 Danes. The sample was drawn to be representative of the adult population as of January 1, 2009, using sampling procedures that are virtually identical to those documented at length in Harrison et al. (2005). We received a random sample of the population aged between 18 and 75, inclusive, from the Danish Civil Registration Office, stratified the sample by geographic area, and sent out 1969 invitations.[10]

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

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

There were 40 discounting choices and 40 risk attitude choices, and each subject had a 10% chance of being paid for one choice on each block. Average payments on the first block were 201 kroner (although some were for deferred receipt) and on the second block the average was 242 kroner, for a combined average of 443 kroner. The exchange rate at the time was close to 5 kroner per U.S. dollar, so earnings averaged \$91 per 2 two-hour session for these two tasks. Subjects were also paid a 300 kroner or 500 kroner fixed show-up fee, plus earnings from subsequent tasks.[11] For payments to be made in the future, the following language explained the procedures:

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

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

4. Econometrics

Our objective is to evaluate alternative discounting functions reviewed in Section 2. The approach we adopt is direct estimation by maximum likelihood of a structural model of the latent choice process in which the core parameters defining risk attitudes and discounting behavior can be estimated. The approach is an extension of the full-information maximum likelihood specification used in Andersen et al. (2008a), of course with modifications for the specification of alternative discounting functions.[12] We review the inferential logic for estimating risk attitudes and discounting behavior and detailed specifications in Appendix E (available online).

5. Results

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

5.1. Initial estimates

Table 1 reports maximum likelihood estimates of the main discounting functions. Underlying each of these sets of estimates are models of the non-linear utility function using a constant relative risk aversion (CRRA) specification, as well as

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

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

[12] Andersen et al. (2008a) estimate a structural version of the dual-self model of "impulse control" developed by Fudenberg and Levine (2006), in which income from the experimental task is integrated into other extra-experimental income and wealth. They assume that income from the tasks that pay out on the day of the experiment is spent in one day, and income from the discount rate tasks that pay out in the future is spent in $\lambda \geq 1$ days. The structural model is estimated from choice behavior in 2003 by a sample of adult Danes. The results show that the fit of the model is maximized when $\lambda = 1$, when income from both tasks are spent over the same period of time, one day. So this extended specification provides direct support for our current specification, from a sample drawn from the same population 6 years earlier.

Table 1
Maximum likelihood estimates of discounting models.

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

the behavioral error parameters[13] and an assumed background consumption value.[14] The point estimate for relative risk aversion is robustly estimated to be 0.65 with a standard error of 0.038, and a 95% confidence interval between 0.58 and 0.73. This is completely consistent with previous findings, and of course implies a concave utility function. To check the validity of the CRRA specification, we followed Harrison et al. (2007) and estimated the more general Expo-Power specification. We could not reject the assumption of CRRA over the domain of prizes, although there was some evidence for very slightly decreasing RRA over that domain.

The estimates in Table 1 show *robust evidence of almost-constant discounting*. There will be statistically significant evidence of non-constant discounting in some specifications, but nothing that is as dramatic in terms of economic significance as the conventional wisdom might suggest. In other specifications, there might be evidence of non-constant discounting in point estimates, but not when one allows for the statistical uncertainty of the estimates. It is not appropriate, of course, to draw inferences from point estimates without considering their statistical precision.

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

The most striking finding from Table 1, for us, is that *there is no quasi-hyperbolic discounting*. The key parameter, β, is not statistically or economically significantly different from 1, and the parameter δ is close to the estimate of δ from the

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

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

Exponential discounting model. The p-value on a test of the hypothesis that $\beta=1$ has value 0.55, although the 95% confidence interval for β is enough to see that it is not significantly different from 1.

We also see from panel C of Table 1 that the rejection of the QH specification is not due to there being a different kind of fixed cost to discounting. We reject the hypothesis from the Fixed Cost discounting models (3a) and (3b) that $\beta<1$, as one might expect from panel B, but we also find no evidence that $b>0$. Furthermore, we cannot reject the joint hypothesis that $\beta=1$ and $b=0$, with a p-value of 0.41.[15] Because $\theta>1$ there is some evidence for hyperbolic discounting, but the statistical significance is very slight. Assuming $\beta=1$ and $b=0$, we estimate θ to be 5.88 with a standard error of 5.33, and one cannot reject the hypothesis with such a standard error that $\theta=1$ (p-value of 0.36). In effect, with $\beta=1$ and $b=0$ this model has collapsed to a Simple Hyperbolic model, and one may as well then estimate it and Generalized Hyperbolic models.

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

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

5.2. Controlling for the treatments

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

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

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

These results suggest that our main conclusion thus far, the lack of support for the QH specification in favor of the exponential model, appears to be robust to controls for the prime suspects in terms of our elicitation procedures. Essentially the same is true for the other specifications.

One of our “treatments,” in a sense, is the elicitation of discount rates over horizons extending from 2 weeks up to one year. To what extent is the lack of evidence for hyperbolicky discounting due to constant-discounting responses to longer horizons swamping non-constant responses to shorter horizons? One could re-weight the data to focus more on the shorter horizons, but a simpler method is to estimate the models with shorter horizons. Focusing on the two shortest horizons of 2 weeks and 1 month, which is roughly 25% of the data due to our deliberate over-sampling design, we do not see any

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

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

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

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

Table 2
Estimates of the effects of treatments.

Parameter	Point estimate	Standard error	p-value	95% Confidence interval	
A. Exponential discounting (LL= −18,459.1; Eq. (1))					
δ Constant	0.117	0.024	< 0.001	0.070	0.164
FED	0.027	0.017	0.126	−0.007	0.061
INFO	−0.036	0.015	0.018	−0.066	−0.006
H_ORDER	−0.037	0.017	0.029	−0.070	−0.004
P_HIGH	−0.001	0.017	0.952	−0.034	0.032
RA_FIRST	0.007	0.020	0.730	−0.032	0.046
FEE_HIGH	−0.014	0.019	0.441	−0.051	0.022
r Constant	0.584	0.062	< 0.001	0.461	0.706
RA_FIRST	0.047	0.066	0.476	−0.083	0.177
FEE_HIGH	0.084	0.064	0.190	−0.042	0.210
B. Quasi-hyperbolic discounting (LL= −18389.4; Eqs. (2a) and (2b))					
β Constant	1.003	0.014	< 0.001	0.976	1.030
INFO	0.013	0.009	0.170	−0.005	0.025
H_ORDER	0.005	0.010	0.607	−0.015	0.025
P_HIGH	−0.005	0.012	0.968	−0.023	0.022
RA_FIRST	−0.023	0.012	0.051	−0.046	0.001
FEE_HIGH	−0.007	0.010	0.498	−0.027	0.013
δ Constant	0.091	0.020	< 0.001	0.051	0.131
FED	0.035	0.017	0.043	0.001	0.069
INFO	−0.021	0.013	0.112	−0.047	0.005
H_ORDER	−0.027	0.016	0.087	−0.058	0.004
P_HIGH	−0.002	0.015	0.873	−0.033	0.028
RA_FIRST	−0.010	0.017	0.578	−0.044	0.024
FEE_HIGH	−0.017	0.016	0.297	−0.049	0.015
r Constant	0.592	0.063	< 0.001	0.469	0.716
RA_FIRST	0.038	0.068	0.571	−0.094	0.171
FEE_HIGH	0.084	0.064	0.192	−0.042	0.210

deviations from constant discounting. There is no statistically significant effect on either parameter of the QH specification, and we reach the same conclusion with the Weibull discounting model.

5.3. Robustness checks

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

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

The second robustness check is to consider mixture models in which the observed choices over time-dated outcomes could be generated by two discounting models rather than one, with some fraction of observed choices accounted for by one model and the remaining choices accounted for by the other model. The mixture model specification jointly estimates the structural parameters of each model as well as the mixing probability between the two of them.[19] As a general matter, we find that all of the major variants "collapse" to being a mixture of two exponential discounting models. In fact, if we estimate that mixture model the smaller discount rate is 6.5% with a weight of 0.79, and the higher discount rate is 11.6% with a weight of 0.21. The p-values on these discount rates estimates are both <0.001.

If we consider a mixture between the exponential and QH discounting models, we find that the QH model collapses to an exponential with $\beta=1$ (p-value of 0.77), and the probability for that model is estimated to be 0.21. Thus it would *appear* that there is some support for the QH specification, until one examines the estimated parameter values for each model: the QH specification is effectively a second exponential specification with a discount rate of 11.5% because the estimate of β is

[19] We note that it is the individual choice that it being classified by the mixture model, not the individual decision-maker. Harrison and Rutström (2009; Section 3.4) discuss why we prefer this approach.

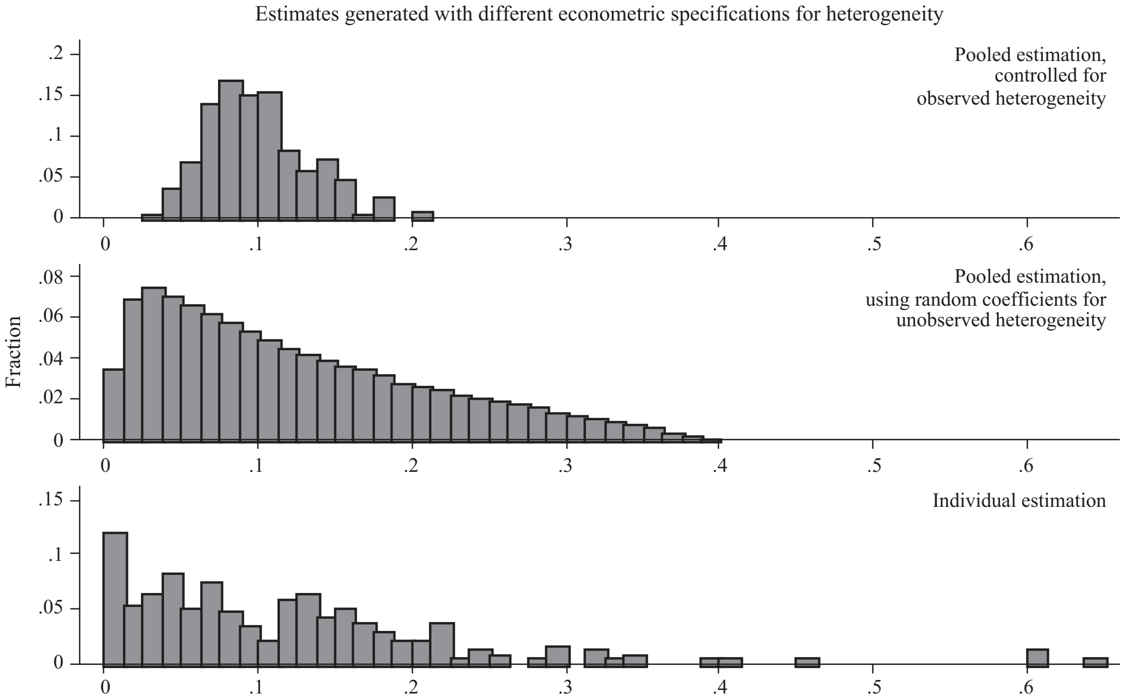

Fig. 1. Estimated distribution of δ assuming the exponential discounting model.

essentially 1. We also consider a mixture between the exponential and Weibull discounting models, and come to the same conclusion.[20] Finally, we consider a mixture between the exponential and simple hyperbolic discounting models, and again come to the same conclusion.[21]

The third robustness check is to allow for observed and unobserved individual heterogeneity in behavior, both in terms of the choices over risk, and inferences about the utility function, and the choices over time-dated outcomes. We first consider observed heterogeneity by means of a set of standard socio-demographic characteristics of individuals affecting each of the structural parameters we estimate. The upper panel in Fig. 1 shows the implied distribution of predicted exponential discount rates from the subjects in our sample: the mean estimate is 10.0% with a standard deviation of 3.2%. The middle panel displays the estimated population distribution of discount rates using random coefficients for unobserved heterogeneity, with a higher mean estimate of 12.5% and a larger standard deviation of 8.9%. Finally, we consider estimation at the level of the individual: the mean is 11.6 and the standard deviation is 10.9%. The distribution of estimated discount rates in the random coefficients model reflects the variation in discount rates at the individual level, which illustrates the complementarity between the two estimation methods. There is considerable variation in estimated discount rates for each of three different estimation methods, with more evidence of variation in the models that allow for unobserved heterogeneity.[22]

[20] In this case the mixture model assigns 81% weight to the exponential specification. The hypothesis test that the Weibull collapses to an exponential with $\acute{s}=1$ has a p-value of 0.083, which is on the borderline of being statistically significant depending on the critical value used. At the 5% significance level one would *not* reject the hypothesis that the Weibull component of the mixture collapses to an exponential. The discount rates implied by the Weibull specification are 7.9% for 1 day, 11.6% for 1 week, 15.6% for 1 month, and 25.8% for 1 year. Although it may seem *a priori* implausible to see increasing discount rates with horizon, the discount factors implied by these discount rates are still decreasing, and the above test of the hypothesis that $\acute{s}=1$, and inspection of Fig. F2 in Appendix F, imply that the 95% confidence interval around these point estimates is quite wide. These findings are consistent with constant discount rates that are relatively imprecisely estimated compared to the Exponential component of the mixture.

[21] In this case the mixture model assigns 80% weight to the exponential specification, with a discount rate of 5.7%. The discount rates for the Hyperbolic specification range between 29% and 25% for horizons of one day and one year. Andersen et al. (2008a) considered an identical mixture model between exponential and simple hyperbolic discounting models for the same population in 2003, and found that 72% of the observations could be characterized by the exponential specification with a discount rate of 6.8%. The remaining 28% of the observations for 2003 were characterized by the simple hyperbolic specification that was statistically significantly different from exponential, and implied discount rates of about 50% for horizons of 3 months, falling to 20% for horizons of 1 year. Our current results, with a 2009 sample and some changes in experimental design and econometrics, are roughly the same in terms of the weights assigned the exponential and hyperbolic specifications, as well as the discount rates implied for the exponential specification. Our current results for the hyperbolic specification exhibit far less of a decline in discount rates with shorter horizons, although they are consistent with implied exponential discount rates for longer horizons such as 1 year.

[22] The models with observable heterogeneity and random coefficients are based on responses from 308 subjects. For comparability, we removed any subjects that always chose sooner options or always chose later options, since we could not estimate at the individual level for them. We solved the model at the individual level for 238 of these subjects.

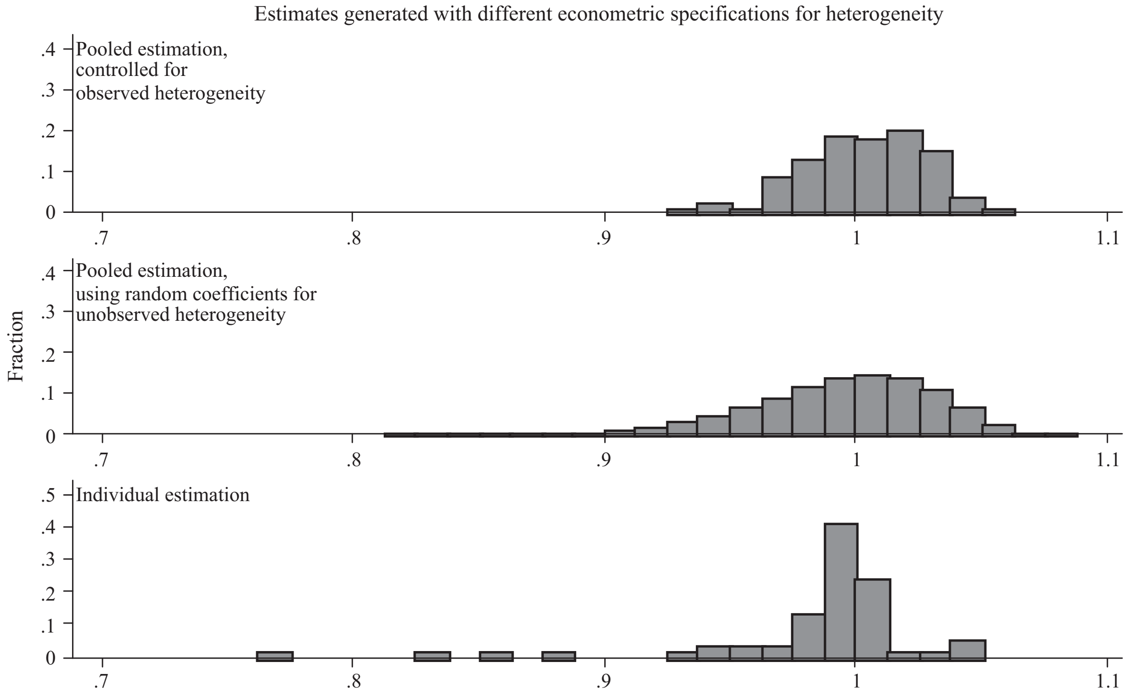

Fig. 2. Estimated distributions of β assuming the quasi-hyberbolic discounting model.

The estimated distributions of the β parameter in the QH model are displayed in Fig. 2. We observe, in the upper panel, that the estimated β coefficient has a mean of 1.00 with a standard deviation of 0.02 when we control for observed heterogeneity. Allowing for unobserved heterogeneity with random coefficients generates a distribution with a mean of 0.99 and standard deviation of 0.03, and undertaking individual estimation generates a distribution with a mean of 1.00 and standard deviation of 0.04. Finally, Fig. 3 shows the estimated distributions of δ in the QH model for different econometric specifications of individual heterogeneity.[23] The results are qualitatively the same as for the exponential model in Fig. 1: the mean (standard deviation) of the distribution is 10% (3%), 11% (9%), and 9% (10%) for the top, middle and bottom panels. None of the three approaches to modeling individual heterogeneity changes our basic conclusions about the best discounting model to characterize these data.

The fourth robustness check is to see if the absence of evidence for hyperbolic discounting is due to the theoretically motivated use of a concave utility function when inferring discount rates from observed choices over time-dated amounts of money. Since a concave utility function significantly lowers the level of discounting inferred, perhaps that means that there is simply less “headroom” for discount rates to be hyperbolic. This is easy to check by assuming a linear utility function for each of our specifications. We certainly infer higher discount rates, but in no case do we observe any statistically significant decline in discount rates with horizon.

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

6. Connection to previous literature

Our results were a surprise to us, and the robustness checks reported above did not lead us to qualify that reaction. We fully expected to see much more “hyperbolicky” behavior when we removed the front end delay, and particularly when

[23] The models with observable heterogeneity and random coefficients are based on responses from 198 subjects in the no-FED treatment. Apart from dropping subjects in the FED treatment, for comparability we also dropped subjects that always chose sooner options or always chose later options. We solved the model at the individual level for 87 of those subjects.

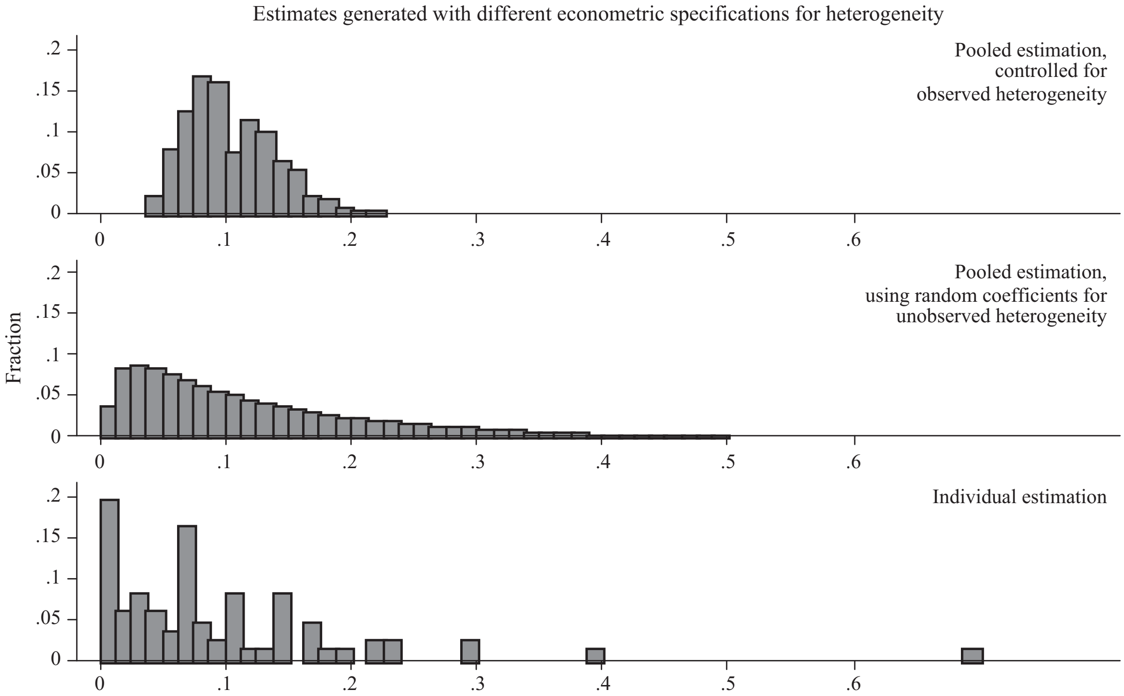

Fig. 3. Estimated distributions of δ assuming the quasi-hyberbolic discounting model.

that was interacted with not providing the implied interest rates of each choice. We were not wedded to one hyperbolic specification or the other, and did not expect the exponential model to be completely overwhelmed by the alternatives, but we did expect to see much more non-constant discounting. We therefore re-examined the literature, and tried to draw some inferences about what might explain the apparent differences in results.

6.1. Reconsideration of previous literature

We undertook a re-examination of the previous, comparable literature that has led to the conventional wisdom of significant non-constant discounting. We ignored all hypothetical survey studies, on the grounds that the evidence is overwhelming that there *can* be huge and systematic hypothetical biases. It is simply inefficient to take the evidence from hypothetical survey studies seriously.[24] We also focused on experiments, rather than econometric inferences from naturally occurring data, because those data are easier to interpret and have generated the conventional wisdom.[25] We excluded studies that did not lend themselves to inferring a discount function.[26] Finally, we excluded any study that used procedures that were not incentive-compatible or that involved deception.[27]

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

---

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

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

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

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

Table 3
Review of experimental literature with real incentives.

Study	Sample (size)	Elicitation method	Horizon (s)	Front end delay (s)	Correct for non-linear utility	Models (estimated rates)	Statistical method	Hyperbolicky discounting?
Ainslie and Haendel (1983)	Patients (N=18, 66 choices)	Choice	3 days	0 or 7 days	No	None (n/a)	Counting	Yes
Horowitz (1991)	Students (N=70)	Bidding	64 and 34 days	None	Yes	None (830%, 271%)	Summary statistic	Yes
Kirby and Maraković (1996)	Students (N=621)	Choice	10–70 days	None	No	E, H2 (128% to 1.2E+13)	Non-linear least squares	Yes
Kirby (1997)	Students and others (N=24, 28, 20)	Bidding	1–29 days	None	No	E, H2 (1.6E+04 to 4.3E+10)	Non-linear least squares	Yes
Coller and Williams (1999)	Students (N=199)	Choice	2 months	0 or 1 month	No	E (72.5% with no FED; 30.2% with FED)	Interval regression	Yes (see our Appendix B)
Kirby et al. (1999)		Choice			No			
Anderhub et al. (2001)	Students	Pricing	4, 8 weeks	None	No	E and H (128–1084%)	Non-parametric tests	No (see our Appendix B)
Harrison et al. (2002)	Danish adults (N=268)	Choice	6, 12, 24 and 36 months	1 month	No	E (28.1%)		No
Kirby and Santiesteban (2003)	Students	Bidding	1–43 days	0 (Experiment 1); 1 day (Experiment 2)	"Not really"	E, H2 (11.3– 2877%; 18–71,231%)		
Eckel et al. (2005)	Canadian adults	Choice	2 to 28 days	0, 1 day, 2 days, or 2 weeks	No	E		
Harrison et al. (2005)	Danish adults (N=243)	Choice	1, 4, 6, 12, 18, 24 months	1 month	No	E (23.8%)	Interval regression	No
Andersen et al. (2008a)	Danish adults (N=243)	Choice	1, 4, 6, 12, 18, 24 months	1 month	Yes	E, H3, W (10.1%)	ML structural estimates	No
Engle-Warnick et al. (2009)	Students (N=151)	Bidding	0, 8, 25 weeks	None	Yes	QH (38% and 33%)	Non-linear least squares	No
Andersen et al. (2010)	Students (N=90)	Choice	1, 4 and 6 months	1 month	No	E (27.9%)	Interval regression	No
Dohmen et al. (2010)	German adults (N=500)	Choivr	12 months	None	No	None	Interval regression	
Takeuchi (2011)	Students (N=56)	Bidding	Elicited	None	"Not really"	E (726%)	Non-linear least squares	Yes
Benhabib et al. (2010)	Students (N=27)	Matching	3 days, 1 and 2 weeks, 1,3 and 6 months	None	No	E and FC ($\approx$ 472%)	Non-linear lest squares	Yes
Coller et al. (2012)	Students (N=87)	Choice	1–60 days	None	Yes	E and QH (mixture) ($\geqslant$1000–33%)	ML structural estimates	Yes
Andreoni and Sprenger (2012a)	Students (N=97)	Portfolio allocation	35, 70 and 98 days	0, 7 and 35 days	Yes	E (30% overall; 28% with no FED)	Non-linear least squares	No
Laury et al. (2012)	Students (N=103)	Choice	9 weeks	3 weeks	Yes	E (12.2% and 14.1%)	ML structural estimates	No

Two additional conclusions from the review of the literature are the potential roles of "small stakes" and open-ended, "fill in the blank" elicitation procedures. Evaluating these characteristics of the previous literature is beyond the scope of our study, since they raise a host of behavioral and experimental issues.[28]

6.2. *Experiments with students*

In order to determine if the evidence for non-constant discounting derives from the general focus on students samples, we replicated our field experiments with a student sample in Copenhagen recruited using standard methods.[29] The experimental tasks were identical, to ensure comparability.

Table 4 lists estimates from the student responses of the basic models in Table 1. The risk attitudes of this sample were close to those of the adult Danish population.[30] The results are intriguing, but do not change our basic conclusions. We do observe a slightly higher discount rate with the exponential model: 10.4% compared to 8.9% for the adult population in general. And we obtain point estimates that suggest modest QH discounting ($\beta=0.986$), although the 95% confidence interval spans $\beta=1$ and we can reject the assumption of $\beta=1$ at the $p=0.18$ level for a two-sided test, and hence at the $p=0.09$ level for the appropriate one-sided test. So this is suggestive of QH discounting, but not clear evidence. We find no evidence of Fixed-Cost discounting, and no evidence of simple hyperbolic discounting. We do observe some non-constancy of discount rates with the Weibull discounting specification, although the overall effect of the student sample is not statistically significant, as shown by the p-value of 0.17 on the null hypothesis that the specification is actually exponential.[31]

6.3. *The probability of discounting*

Our literature review deliberately ignored studies that do not consider decisions with real monetary consequences, but there is one treatment that we do want to recognize even if it has only been addressed in studies using hypothetical survey questions: the effect of the discounting tasks being rewarded probabilistically.[32] In our experiments each subject had a 10% chance that one of their discounting tasks would be rewarded.[33]

Keren and Roelofsma (1995) demonstrated, with hypothetical tasks, a behavioral effect of this treatment on discounting behavior, specifically a reduction in the extent of hyperbolicky behavior for shorter horizons. Their first experiment illustrates their findings. Subjects were offered 100 Dutch Guilders now or 110 in 4 weeks. When the payment was not probabilistic, 82% of 60 subjects chose the sooner option. But when the payment would occur with a probability of 0.9, only 54% of 70 subjects chose the sooner option, and that declined to only 39% of 100 subjects with a probability of 0.5. With a front end delay of 26 weeks, the same subjects chose the sooner option in 37%, 25% and 33% of the choices, respectively, suggesting that reducing the probability of payment to 0.5 generated results consistent with having a front end delay.

Weber and Chapman (2005) were unable to replicate these findings.[34] They used 446 students in an introductory psychology class in a between-subjects replication of the experiment described above, but with U.S. dollar amounts instead of Dutch Guilders. With no front end delay the fraction choosing the sooner option was 61% (of 113) with a probability of 1 of

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

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

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

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

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

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

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

Table 4
Estimates of discounting models with student sample.

Parameter	Point estimate	Standard error	*P*-Value	95% Confidence interval	
A. Exponential discounting (LL = −3441.9; Eq. (1))					
δ	0.104	0.014	< 0.001	0.076	0.131
r	0.534	0.056	< 0.001	0.424	0.644
B. Quasi-hyperbolic discounting (LL = −3427.9; Eqs. (2a) and (2b))					
β	0.986	0.010	< 0.001	0.966	1.006
δ	0.072	0.012	< 0.001	0.048	0.096
r	0.532	0.056	< 0.001	0.423	0.642
H_0: $\beta=1$, p-value = 0.18					
C. Fixed cost hyperbolic discounting (LL = −3418.9; Eqs. (3a) and (3b))					
θ	10.167	26.28	0.699	−41.341	61.674
β	0.966	0.031	< 0.001	0.905	1.026
δ	0.123	0.168	0.448	−0.202	0.458
b	−0.060	0.063	0.340	−0.185	0.064
r	0.529	0.056	< 0.001	0.421	0.638
H_0: $\beta=1$, p-value = 0.27; H_0: $\beta=1$ and $b=0$, p-value = 0.53					
D. Simple hyperbolic discounting (LL = −3440.6; Eq. (4))					
K	0.103	0.014	< 0.001	0.076	0.131
r	0.534	0.056	< 0.001	0.424	0.643
E. Generalized hyperbolic (LL = −3427.9; Eq. (5))					
α	5.540	7.530	0.462	−9.219	20.299
β	0.261	0.191	0.173	−0.114	0.636
r	0.532	0.056	< 0.001	0.423	0.642
F. Weibull discounting (LL = −3425.9; Eq. (6))					
$\acute{r}$	0.094	0.012	< 0.001	0.070	0.118
$\acute{s}$	1.594	0.428	< 0.001	0.758	2.430
r	0.532	0.056	< 0.001	0.423	0.642
H_0: $\acute{s}=1$, p-value = 0.17					

payment, and 70% (of 111) with a probability of 0.5. Adding a front end delay generated comparable choices of 46% (of 109) and 51% (of 113), implying no significant effects of having probabilistic payments. Although one hesitates to pursue design differences with non-salient tasks, it is worth noting that the Keren and Roelofsma (1995) subjects were compensated for attending the session.

We consider the effect of probabilistic discounting by undertaking experiments in which we vary the exogenous probability of payment. Specifically, we conduct experiments with 28 subjects from the greater Copenhagen area in which we vary the probability of payment for the discounting task from 10% to 100%, and see if there is a difference in behavior. Of course, increasing the probability means that we need to account for the scale effects on expected rewards. In Keren and Roelofsma (1995) the stakes were kept the same, so there may be a confound of a scale effect. For example, their subjects might have been close to risk-neutral for lower stakes (hence implying higher discount rates when the stakes were to paid with some probability less than 1) and risk averse for higher stakes (hence implying lower discount rates when the probability of payment was closer to 1). This pattern of risk aversion is found in many laboratory settings: for example, see Holt and Laury (2002, 2005) and Harrison et al. (2005). We therefore maintain the stakes at their original levels, despite the cost of the experiments, and allow for varying risk aversion with stakes.

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

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

The new text was simply this:

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

(*footnote continued*)
original, involving the use of dubious indifference-point elicitation procedures. So although it changed the *design* from the original, it did replicate the *finding* from the original.

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

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

For example, for the exponential model we first estimate δ_0 and δ_1 in $\delta=\delta_0+\delta_1 \times C$, where C is a binary indicator variable for the 100% certain responses. There is no statistically significant effect on the discounting parameter(s) for the exponential, simple hyperbolic, fixed cost hyperbolic, generalized hyperbolic, and Weibull models.

For the QH discounting model there is an effect on the all-important β parameter.[35] It is 0.025 lower with the 100% payment treatment, and this effect has a p-value of 0.074. Estimating the QH model with the new sample we estimate β to be 0.983 with a standard error of 0.013, and δ to be 0.041 with a standard error of 0.024. The hypothesis that $\beta=1$ has a p-value of 0.185, so this is not statistically significant evidence in favor of the QH model.

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

7. Open issues

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

First, we need to allow for alternative models of decision-making over risk for some decision-makers in some settings. The identification of non-standard models of risk preferences, such as RDU or Cumulative Prospect Theory, demands careful attention to the tasks given to subjects, and is not something that we believe can be safely "folded in" with some other task, as proposed by Andreoni and Sprenger (2012a). In a similar vein, interacting risk and time preferences, say by offering subjects a choice over time-dated lotteries, may raise deep confounds if one insists on standard, additive intertemporal utility functions. The claim that "risk preferences are not time preferences" of Andreoni and Sprenger (2012b) can be viewed as an illustration of that confound at work (see Andersen et al., 2011b).

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

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

Related to this point, one intriguing hypothesis behind our finding that Danes tend to discount in an exponential manner could be that our experiment served as an artefactual commitment device. The idea is that our highly structured experiments, and the formal manner in which we explained the credibility of our making payments on specified future dates, might appeal to individuals that have difficulties making consistent intertemporal plans in terms of other choices they make. This is a hypothesis about the external validity of our findings, and the extent to which they might transfer to other, less-structured intertemporal choices. We find this hypothesis attractive *a priori*, and just note that such control and structure may be the price one pays for internal validity in claiming to have measured time preferences. Building a bridge between that internal validity and a wider class of field choices is an important challenge for future experiments (Harrison and List, 2004).

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

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

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

8. Conclusions

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

How do we reconcile this striking finding with the received wisdom? We see nothing in our experimental procedures which might bias behavior, and that deviates in any novel manner from the types of procedures used in the past. We avoid eliciting present values in an open-ended manner, because we are suspicious of the behavioral accuracy of those responses.[36] We test for the effect of providing information on the implied interest rates we offered. We use displays of the tasks that make them relatively transparent in terms of the choice alternatives, rather than rely entirely on the ability of subjects to read numbers and words. And, obviously, we pay the subjects in a salient manner.

Our basic econometric procedures are familiar from the binary choice literature, and have a long tradition in experimental economics (e.g., Camerer and Ho, 1994; Hey and Orme, 1994). Our application of them uses parametric methods, but we are clearly flexible in terms of the discounting functions we examine. The notion of joint estimation of utility functions and discounting functions is driven by theory, and implies nothing fundamental from an econometric perspective. The application of mixture specifications to explore the robustness of our basic results is, similarly, not fundamentally novel in terms of method.

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

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

Appendix A. Supplementary materials

Supplementary data associated with this article can be found in the online version at http://dx.doi.org/10.1016/j.euroecorev.2014.06.009.

References

Ainslie, George, 1975. Specious reward: a behavioral theory of impulsiveness and impulse control. Psychol. Bull. 92 (4), 463–496.

Ainslie, George, 1992. Picoeconomics. Cambridge University Press, New York.

Ainslie, George, Haendel, V., 1983. The motives of the will. In: Gottheil, E., Druley, K., Skodola, T., Waxman, H. (Eds.), Etiology Aspects of Alcohol and Drug Abuse, Charles C. Thomas, Springfield, IL.

Ahlbrecht, Martin, Weber, Martin, 1997. An empirical study on intertemporal decision making under risk. Manage. Sci. 43 (6), 813–826.

Anderhub, Vital, Güth, Werner, Gneezy, Uri, Sonsino, Doron, 2001. On the Interaction of risk and time preference: an experimental study. Ger. Econ. Rev. 2 (3), 239–253.

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

Andersen, Steffen, Harrison, Glenn W., Lau, Morten I., Rutström, E. Elisabet, 2008a. Eliciting risk and time preferences. Econometrica 76 (3), 583–619.

Andersen, Steffen, Harrison, Glenn W., Lau, Morten I., Rutström, E. Elisabet, 2008b. Lost in state space: are preferences stable? Int. Econ. Rev. 49 (3), 1091–1112.

Andersen, Steffen, Harrison, Glenn W., Lau, Morten I., Rutström, E. Elisabet, 2010. Preference heterogeneity in experiments: comparing the field and laboratory. J. Econ. Behav. Organ. 73, 209–224.

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

Andersen, Steffen, Harrison, Glenn W., Lau, Morten I., Rutström, E. Elisabet, 2013. Discounting behavior and the magnitude effect. Economica 80, 670–697.

Andreoni, James, Sprenger, Charles, 2012a. Estimating time preferences from convex budgets. Am. Econ. Rev. 102 (7), 3333–3356.

Andreoni, James, Sprenger, Charles, 2012b. Risk preferences are not time preferences. Am. Econ. Rev. 102 (7), 3357–3376.

Benhabib, Jess, Bisin, Alberto, Schotter, Andrew, 2010. Present-bias, quasi-hyperbolic discounting, and fixed costs. Games Econ. Behav. 69 (2), 205–223.

Camerer, Colin, Ho, Teck-Hua, 1994. Violations of the betweenness axiom and nonlinearity in probability. J. Risk Uncertain. 8, 167–196.

Coller, Maribeth, Harrison, Glenn W., Rutström, E. Elisabet, 2012. Latent process heterogeneity in discounting behavior. Oxf. Econ. Pap. 64, 375–391.

Coller, Maribeth, Williams, Melonie B., 1999. Eliciting individual discount rates. Exp. Econ. 2, 107–127.

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

Cubitt, Robin P., Read, Daniel, 2007. Can intertemporal choice experiments elicit time preferences for consumption? Exp. Econ. 10, 369–389.

Dohmen, Thomas, Falk, Armin, Huffman, David, Sunde, Uwe, 2010. Are risk aversion and impatience related to cognitive ability? Am. Econ. Rev. 100 (3), 1238–1260.

Ebert, Jane E.J., Prelec, Drazen, 2007. The fragility of time: time-insensitivity and valuation of the near and far future. Manage. Sci. 53 (9), 1423–1438.

Eckel, Catherine, Johnson, Cathleen, Montmarquette, Claude, 2005. Savings decisions of the working poor: short- and long-term horizons (Research in Experimental Economics). In: Carpenter, J., Harrison, G.W., List, J.A. (Eds.), Field Experiments in Economics, vol. 10, JAI Press, Greenwich, CT.

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

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

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

Frederick, Shane, Loewenstein, George, O'Donoghue, Ted, 2002. Time discounting and time preference: a critical review. J. Econ. Lit. 40, 351–401.

Fudenberg, Drew, Levine, David K., December 2006. A dual-self model of impulse control. Am. Econ. Rev. 96 (5), 1449–1476.

Halevy, Yoram, 2008. Strotz meets allais: diminishing impatience and the certainty effect. Am. Econ. Rev. 98 (3), 1145–1162.

Harrison, Glenn W., 1992. Theory and misbehavior of first-price auctions: reply. Am. Econ. Rev. 82, 1426–1443.

Harrison, Glenn W., 2005. Field experiments and control (Research in Experimental Economics). In: Carpenter, J., Harrison, G.W., List, J.A. (Eds.), Field Experiments in Economics, vol. 10, JAI Press, Greenwich, CT.

Harrison, Glenn W., 2006. Hypothetical bias over uncertain outcomes. In: List, J.A. (Ed.), Using Experimental Methods in Environmental and Resource Economics, Elgar, Northampton, MA.

Harrison, Glenn W., Johnson, Eric, McInnes, Melayne M., Rutström, E. Elisabet, 2005. Risk aversion and incentive effects: comment. Am. Econ. Rev. 95 (3), 897–901.

Harrison, Glenn W., Lau, Morten, Rutström, Elisabet, 2007. Estimating risk attitudes in Denmark: a field experiment. Scand. J. Econ. 109 (2), 341–368.

Harrison, Glenn W., Lau, Morten I., Rutström, E. Elisabet, Sullivan, Melonie B., 2005. Eliciting risk and time preferences using field experiments: some methodological issues (Research in Experimental Economics). In: Carpenter, J., Harrison, G.W., List, J.A. (Eds.), Field Experiments in Economics, vol. 10, JAI Press, Greenwich, CT.

Harrison, Glenn W., Lau, Morten I., Williams, Melonie B., 2002. Estimating individual discount rates for Denmark: a field experiment. Am. Econ. Rev. 92 (5), 1606–1617.

Harrison, Glenn W., List, John A., 2004. Field experiments. J. Econ. Lit. 42 (4), 1013–1059.

Harrison, Glenn W., Rutström, E. Elisabet, 2008a. Risk aversion in the laboratory (Research in Experimental Economics). In: Cox, J.C., Harrison, G.W. (Eds.), Risk Aversion in Experiments, vol. 12, Emerald, Bingley, UK.

Harrison, Glenn W., Rutström, E. Elisabet, 2008b. Experimental evidence on the existence of hypothetical bias in value elicitation experiments. In: Plott, C.R., Smith, V.L. (Eds.), Handbook of Experimental Economics Results, Elsevier Press, New York.

Harrison, Glenn W., Rutström, E. Elisabet, 2009. Expected utility and prospect theory: one wedding and a decent funeral. Expl. Econ. 12 (2), 133–158.

Harstad, Ronald M., 2000. Dominant strategy adoption and bidders' experience with pricing rules. Exp. Econ. 3 (3), 261–280.

Hey, John D., Orme, Chris, 1994. Investigating generalizations of expected utility theory using experimental data. Econometrica 62 (6), 1291–1326.

Holden, Stein T., Shiferaw, Bekele, Wik, Mette, 1998. Poverty, market imperfections and time preferences of relevance for environmental policy? Environ. Dev. Econ. 3, 105–130.

Holt, Charles A., Laury, Susan K., 2002. Risk aversion and incentive effects. Am. Econ. Rev. 92 (5), 1644–1655.

Holt, Charles A., Laury, Susan K., 2005. Risk aversion and incentive effects: new data without order effects. Am. Econ. Rev. 95 (3), 902–912.

Horowitz, John K., 1991. Discounting money payoffs: an experimental analysis. In: Kaish, S., Gilad, B. (Eds.), Handbook of Behavioral Economics, vol. 2B. , JAI Press, Greenwich.

Keller, L. Robin, Strazzera, Elisabetta, 2002. Examining predictive accuracy among discounting models. J. Risk Uncertain. 24 (2), 143–160.

Keren, Gideon, Roelofsma, Peter, 1995. Immediacy and certainty in intertemporal choice. Organ. Behav. Hum. Decis. Process. 63 (3), 287–297.

Kirby, Kris N., 1997. Bidding on the future: evidence against normative discounting of delayed rewards. J. Exp. Psychol. Gen. 126, 54–70.

Kirby, Kris N., Guastello, Barbarose, 2001. Making choices in anticipation of similar future choices can increase self-control. J. Exp. Psychol. Appl. 7 (2), 154–164.

Kirby, Kris N., Maraković, Nino N., 1995. Modeling myopic decisions: evidence of hyperbolic delay-discounting within subjects and amounts. Organ. Behav. Hum. Decis. Process. 64, 22–30.

Kirby, Kris N., Maraković, Nino N., 1996. Delay-discounting probabilistic rewards: rates decrease as amounts increase. Psychon. Bull. Rev. 3 (1), 100–104.

Kirby, Kris N., Petry, Nancy M., Bickel, Warren, 1999. Heroin addicts have higher discount rates for delayed rewards than non-drug-using controls. J. Exp. Psychol. Gen. 128 (1), 78–87.

Kirby, Kris N., Santiesteban, Mariana, 2003. Concave utility, transactions costs, and risk in measuring discounting of delayed rewards. J. Exp. Psycho. Learn. Mem. Cognit. 29 (1), 66–79.

Laibson, David, 1997. Golden eggs and hyperbolic discounting. Q. J. Econ. 112 (2), 443–477.

Laury, Susan K., McInnes, Melayne Morgan, Swarthout, J. Todd, 2012. Avoiding the curves: direct elicitation of time preferences. J. Risk Uncertain. 44, 181–217.

Loewenstein, George, O'Donoghue, Ted, Rabin, Matthew, 2003. Projection bias in predicting future utility. Q. J. Econ. 118 (4), 1209–1248.

Loewenstein, George, Prelec, Drazen, 1992. Anomalies in intertemporal choice: evidence and interpretation. Q. J. Econ. 107, 573–592.

Mazur, James E., 1984. Tests of an equivalence rule for fixed and variable reinforcer delays. J. Exp. Psychol. Anim. Behav. Process. 10 (4), 426–437.

Phelps, Edmund S., Pollak, Robert A., 1968. On second-best national saving and game-equilibrium growth. Rev. Econ. Stud. 35, 185–199.

Prelec, Drazen, 2004. Decreasing impatience: a criterion for non-stationary time preference and 'hyperbolic' discounting. Scand. J. Econ. 106 (3), 511–532.

Read, Daniel, 2001. Is time-discounting hyperbolic or subadditive? J. Risk Uncertain. 23 (1), 5–32.

Read, Daniel, Frederick, Shane, Orsel, Burcu, Rahman, Juwaria, 2005. Four score and seven years from now: the date/delay effect in temporal discounting. Manage. Sci. 51 (9), 1326–1335.

Read, Daniel, Roelofsma, Peter H.M.P., 2003. Subadditive versus hyperbolic discounting: a comparison of choice and matching. Organ. Behav. Hum. Decis. Process. 91, 140–153.

Rutström, E. Elisabet, 1998. Home-grown values and the design of incentive compatible auctions. Int. J. Game Theory 27 (3), 427–441.

Takeuchi, Kan, 2011. Non-parametric test of time consistency: present bias and future bias. Games Econ. Behav. 71 (2), 456–478.

Warner, John T., Pleeter, Saul, 2001. The personal discount rate: evidence from military downsizing programs. Am. Econ. Rev. 91 (1), 33–53.

Weber, Bethany, Chapman, Gretchen B., 2005. The combined effects of risk and time on choice: does uncertainty eliminate the immediacy effect? does delay eliminate the certainty effect? Organ. Behav. Hum. Decis. Process. 96, 104–118.