

Compensated Discount Functions - An Experiment on Integrating Rewards with Expected Income*

Attila Ambrus[†]

Duke University, Department of Economics

Tinna Laufey Ásgeirsdóttir[‡]

University of Iceland, Department of Economics

Jawwad Noor[§]

Boston University, Department of Economics

László Sándor[¶]

Harvard University, Department of Economics

April 2014

*We thank Parag Pathak and Georg Weizsäcker for participating in the earlier stages of the project, especially for Georg Weizsäcker's input in designing and running the experiment. We also thank the following research assistants who helped us at various stages of the project: Adalheidur Osk Gudlaugsdottir, Hjorleifur Palsson, Thorhildur Olafsdottir, Kristofer Gunnlaugsson, Kristin Eiriksdottir, Kari S. Fridriksson, Katrin Gunnarsdottir, Valur Thrainsson, Jonas Orn Helgason and Dagny Osk Ragnarsdottir. Financial support from the National Science Foundation (grant SES-0822927), the University of Iceland Research Fund and The Icelandic Research Fund (IRF grant number 080604043) is acknowledged. Noor thanks the MIT Department of Economics for their hospitality, where part of this research was completed.

[†]E-mail: aa231@duke.edu

[‡]E-mail: ta@hi.is

[§]E-mail: jnoor@bu.edu

[¶]E-mail: sandor@fas.harvard.edu

Abstract

This paper examines the open empirical question of whether subjects, when choosing among rewards at different points of time, integrate the latter with their income from other sources, which would cause their background marginal utility for money to affect their preferences over rewards. Moreover, we seek to measure time preferences after compensating for possible income effects. Besides eliciting subjects' preference between standard delayed rewards, the experimental design also elicited their preferences over delayed rewards that are received only if the subject's income remains approximately constant. These preferences along with elicited subjective probabilities of satisfying the condition make the correction possible. We conducted the experiments in Iceland, where prompt availability of income tax returns enabled us to condition delayed rewards on income realizations.

We find that background income affects preferences over unconditional delayed rewards. While most people exhibited present bias when comparing unconditional delayed rewards, subjects with stable income did not. The results are similar for the entire sample once we correct subjects' discount functions for income effects. This suggests that income expectations have an effect on choices between future rewards, and that this may account for some of the present-bias observed in experiments.

Keywords: time preferences, hyperbolic discounting, income expectations, rewards conditional on income realization

JEL Classification: C93, D03, D11, D90

1 Introduction

In order to make inferences from their trade-offs between delayed monetary rewards, experimental studies exploring the nature of time preferences typically presume that a subject's marginal utility for money is constant across time. However, several theoretical papers note that subjects may integrate these rewards with their baseline consumption levels (Olson and Bailey 1981, Rubinstein 2002, Frederick et al. 2002, Noor 2009 and Gerber and Rohde 2010). In this case, anticipated changes in marginal utility for money would influence their trade-offs between delayed rewards. This is also related to the recent experimental and theoretical literature that accounts for an unavoidably uncertain future as a contrast to a certain present.¹ On the other hand, it may be argued that integration with baseline consumption plans might be too difficult for people to do,² or that small rewards are often viewed as windfalls under different mental accounting and enjoyed separately from base consumption. This important open empirical question of whether base income matters for intertemporal choices is the subject of this paper.³ We show how to compensate for possible income effects when measuring discount functions.

To test for the effects of baseline income on evaluation of delayed rewards, and to measure "pure" time discounting (in the sense of being independent of expected marginal utility of consumption), we conducted a lab experiment. We did this in Reykjavik, Iceland, using a random sample of individuals from

¹See Weber and Chapman (2005), Fernandez-Villaverde and Mukhijit (2011), Baucells and Haukamp (2012), Andreoni and Sprenger (2012) and Halevy (2012).

²See some arguments along this line in pages 356-357 of Frederick et al. (2002).

³As in the literature, we will assume that rewards are consumed in the period received, so that "integration with income" just means that "the marginal utility of money matters for the ranking of rewards". A separate strand of the literature focuses on timing a consumption stream, or primary rewards (McClure et al., 2007), as the value of timed monetary rewards also change with liquidity constraints and other financial frictions, as well as *in vivo* transfers.

the census, conditional on them living in post codes not too far from the lab.⁴ The possibility of prompt access to individual tax information in Iceland allowed us to both observe subjects' income for the two years preceding and following the experiment, and to use a novel experimental design in which we offered future rewards *conditional* on particular income realizations.

In the first part of the experiment, we used a standard design of asking participants to choose between unconditional present rewards and unconditional rewards received one or two years later.⁵ The elicitation was done through a series of binary choice questions, presented as a standard multiple price list. From this we derive the *uncompensated discount function*: for instance, if a subject is indifferent between \$100 in the current period and \$150 in t years, we obtain $D_u(t) = \frac{100}{150}$. This discount function is uncompensated in the sense that expected income changes create a wedge between the indifference point and the subject's true underlying discount function.

In the second part of the design, we asked participants to choose between (a) unconditional present rewards and (b) rewards received one or two years later conditional on the subject's income staying "approximately" constant.⁶ The idea is that these rewards get integrated with the same baseline income as present rewards and thus are assessed under the same marginal utility

⁴This feature of taking lab experiments to field subjects is similar to Andersen et al. (2008). Note that such subject pools very likely include more people with stable income than the canonical lab subjects of students. As consumptions of the latter are highly volatile (and steeply rising in a matter of years), any confound of income trends we document in our subsample is likely to be even more powerful for subject pools biased towards students.

⁵We used a front-end delay of one week to put both options on an even footing with respect to transaction costs, immediacy, or trust in the experimenters. What we label here as present rewards were specified to be paid a week after the experiment. Similarly, rewards labeled one and two years later were paid one year plus one week, and two years plus one week later.

⁶More precisely, we required that the after-tax inflation-adjusted income of the participant stayed within 4% of base income, for both the month and the year preceding the moment in question. See our online appendix for the translation of our actual questionnaire.

of money. The evaluation of a conditional delayed reward depends on the utility from receiving the reward, the degree to which it is discounted due to temporal delay, and the subject’s *beliefs* about the likelihood that the conditioning event will be satisfied. Eliciting the second element – the true discount function – is our objective. By assuming that the rewards are small relative to background income, we assume approximate linearity of utility from the rewards. We elicit beliefs in an incentivized fashion as follows. First, we sought a good whose utility to the subject is plausibly independent of the marginal utility of money – we consider charity payments made on behalf of the subject to possess this property. Next, we ask subjects to compare (a’) a payment h made to a charity of the subject’s choice at time t with probability α and (b’) a payment h to the same charity at time t conditional on the subject’s income staying roughly the same. By observing preferences for various α , we uncover the desired subjective belief. With this in hand, we can then derive the underlying discount function (see Section 2 for details), which we label as the *compensated discount function*, to contrast it with the original uncompensated discount functions that do not take income expectations into account.

For the analysis, we assume that subjects face borrowing and lending constraints (frictions) that induce them to consume small rewards at the time when they are received. This assumption is implicitly made in most experimental investigations of time preferences. Indeed, if subjects could freely transfer monetary amounts across time periods then they would only care about the discounted present value of their earnings. In that case, even non-exponential discounters’ choices would look consistent with exponential discounting with the interest rate as the discount rate. The fact that evidence for non-exponential discounting exists implies that the consumption of rewards is not smoothed over time. Our results in this paper suggest in turn that subjects’ marginal utility for money tracks their income, and therefore that subjects are not able to smooth background consumption either.

We also assume that the money amounts we offer subjects are “small” relative to their background consumption, in which case money adds to the agent’s background utility approximately linearly. The rewards we offered subjects were equivalent to approximately \$165 (in 2010) and we assume this to be “small” relative to a period’s income, which in our analysis was 12 months.

Our first set of results only uses the traditional experimental questions that use unconditional future rewards. Consistent with most existing studies, we find that in the full sample the estimated discount factor between one and two years from the experiment is significantly higher than the estimated discount factor between the time of the experiment and one year later.⁷ In the quasi-hyperbolic “ β - δ ” framework, the estimated present-bias parameter β is 0.89 and significantly different from 1, while the estimated long-run discount factor δ is 0.92. We get a very different picture though when we look at subjects with stable incomes. In particular, when we restrict attention to subjects whose real annual income remained within a 10% range of current income in both of the two years following the experiment (31 of the 116 subjects), the discount factor between the first and second year after the experiment is almost identical to the discount factor between the present and one year after the experiment, and the estimated β is 0.97 and not significantly different from 1 (while the estimated δ for this group is almost exactly the same as for the whole subject population). This means that those subjects whose real income remained stable after the experiment made choices consistent with exponential discounting, while other subjects

⁷Many experiments describe discount functions as hyperbolic, implying a present bias. See Frederick et al. (2002) for a review, and more recent evidence by Benhabib et al. (2012). Pender (1996) finds similar results in a field experiment using rice instead of money for rewards. Other papers find no evidence for hyperbolic discounting: see for example Harrison et al. (2002), Andersen et al. (2008) and Andreoni and Sprenger (2012). The experimental finding of present bias generated much theoretical work such as Laibson (1997), O’Donoghue and Rabin (1999), Gul and Pesendorfer (2001) and Fudenberg and Levine (2006), and led to much applied work.

on average made choices that revealed significant present bias.

We get similar results when, instead of realized income, we use subjects' expectations, elicited in an unincentivized survey, at the time of the experiment to identify those who expect stable income in the two years following the experiment. For example, for those subjects who expect to stay in their current job with more than 80% probability for the two years following the experiment, the estimated discount factor between one and two years from the experiment is not significantly different than the discount factor between the time of the experiment and one year later, and the estimated β is 0.94.

We also investigate subjects whose income increases significantly in the year following the experiment, but stays relatively stable afterwards.⁸ Given that the expected marginal utility of extra income for these subjects (as long as they could foresee the said income pattern) is lower in both future years than at the time of the experiment, but about the same magnitude in the two future years, we expect a more significant present bias in these cases. In line with the theoretical predictions, the estimated β decreases to 0.79.⁹

To further examine the relationship between income stability and exponential discounting, in the second half of the analysis we investigated subjects' choices regarding conditional rewards, and derived their compensated discount functions. We had two concerns that needed to be addressed for our data analysis. First, our derivation of the compensated discount functions assumed that subjects followed the discounted expected utility model, and to the extent that this misspecifies some of the subjects' model, we may not get meaningful compensated discount functions for all subjects. Second, the conditional questions are cognitively more demanding than the simpler traditional unconditional questions, and therefore we had to expect that some

⁸In the standard model, a one-shot future reduction (resp. increase) in marginal utility of money predicts present (resp. future) bias in unconditional choices, confounding the true curvature of the utility function with uncertainty and trends in consumption.

⁹There were few subjects with a substantial reduction in income to allow a test of the corresponding prediction of higher estimated β .

subjects’ responses may not be meaningful (for this purpose we put in much effort to streamline the experimental questionnaire and provide a detailed instruction session). Consequently, we conducted our analysis based not just on the whole sample, but also on a restricted subsample of subjects whose answers conveyed “sensible” compensated discount factors. Without imposing arbitrary restrictions on the data based either on the results or some covariates, we fit a finite mixture of bivariate normal distributions on the one-period discount factors and the ones between years 1 and 2 maximizing the Bayesian Information Criterion, with discrete classification into clusters. The number of mixtures fitted was chosen by the same algorithm and the same criterion, also allowing for classification of observations as outliers (noise). Out of the resulting three clusters, two show compensated discount rates in a reasonable range, consisting of 68.1% of the sample (79 observations overall). The likelihood of falling into either of these two clusters is positively correlated with the subject’s education level.

We ran multiple tests of the compensated discount rates being equal to the uncompensated ones (using their difference), or either being 1 (the benchmark of treating future and present equivalently). For β corresponding to present-bias in the quasi-hyperbolic discounting framework, we also tested for differences between the compensated and the uncompensated cases (using their difference), as well as either measure being different from 1. For all these, we employed two-sided t -tests using standard errors robust to heteroskedasticity, and also non-parametric Wilcoxon signed rank tests.

Our first finding is that compensated discount factors are higher than uncompensated ones, for the entire two-year time period covered in our experiment. This suggests that experiments that do not take into account expectations regarding future income overestimate subjects’ true impatience. Second, we find that compensated discount functions are significantly less present biased than uncompensated ones ($p = 0.24$ for the arithmetic difference being 0, but $p = 0.033$ with the signed rank test). The estimated β

parameter, when using compensated discount factors, is higher than when using uncompensated discount factors, and not significantly different from 1, while it is significantly less than one when using uncompensated discount factors.

To summarize, our two strands of analysis point in the same direction. Both suggest that people integrate promised experimental rewards with their underlying expected income, and that their inherent impatience is less than what traditional experiments eliciting time preferences tend to find. When differences in expected future incomes are compensated for, the implied discount function becomes less hyperbolic, and closer to exponential.

Our paper complements several recent papers that find a relationship between income shocks and time preferences, in different contexts: Harrison et al. (2005), Epper (2010), Krupka and Stephens (2013), Carvalho et al. (2014) and Dean and Sautmann (2014).¹⁰ These papers find a relationship between subjects' income situation changes and their uncompensated discount functions.¹¹ Our paper additionally elicits beliefs about future income, and directly observes the background income process of the subjects around the time of the experiment, giving us a picture of both anticipated and realized income. Moreover, we contribute to the literature by constructing an experimental design to estimate the compensated discount function. That is, beside noting the confounding of pure time preference and background income effects, we also show how to isolate and measure the former.

The rest of the paper is structured as follows. Section 2 presents the theory of rewards integrated with other income, including our formula for recovering primitives from elicited choices and subjective probabilities. Section 3 details the experimental design, section 4 the conduct. Section 5 gives

¹⁰On the other hand, Meier and Sprenger (2010) and Gine et al. (2013) find no correlation between financial shocks and time preferences.

¹¹Epper (2010) additionally focuses on the role of anticipated income situation changes. He finds a significant difference between time preferences of college juniors and seniors, and connects it to the expected time before graduation.

the empirical results, while Section 6 concludes.

2 Theoretical background

Here we provide theoretical foundations for the experiments investigating conditional discount factors.

2.1 Overview

Suppose that subjects evaluate consumption using an expected discounted utility model with (uncertain) background consumption b_t . As we compute below, these subjects will evaluate an unconditional reward m received at time t by $D(t) \cdot E[u(b_t+m) - u(b_t)]$, that is, the discounted increase in expected utility due to the reward. Moreover, assuming that m is small relative to b_t , we can exploit the approximation that $u(b_t+m) - u(b_t) = u'(b_t) \cdot m$. Therefore, the utility of an unconditional reward is

$$D(t) \cdot E(u'(b_t)) \cdot m.$$

Denote current income by b^* , which is of course not uncertain.

Traditional experiments elicit the present amount X_t^u which is just as good as a future unconditional reward, yielding the equality

$$u'(b^*) \cdot X_t^u = D(t) \cdot E(u'(b_t)) \cdot m.$$

Define the *uncompensated discount function* by

$$D_u(t) = \frac{X_t^u}{m}. \tag{1}$$

Seeing that $D_u(t) = D(t) \cdot \frac{E(u'(b_t))}{u'(b^*)}$, it is clear that the uncompensated discount function correctly estimates the true discount function $D(t)$ if and only if $\frac{E(u'(b_t))}{u'(b^*)} = 1$ for all t , which will not hold generically, and is the reason

for the usual assumption of constant marginal utility $u'(b_t)$ across time in experiments that use $D_u(t)$ as an estimate for $D(t)$.

To compensate for income effects (ie, non-constant marginal utility across time), consider payments that are paid only in the “constant income” event that $b_\tau = b^*$ for all $\tau \leq t$. Denote this event by s^t . Then such conditional rewards are evaluated by the discounted increase in expected utility:

$$D(t) \cdot p(s^t) \cdot u'(b^*) \cdot m.$$

So, if an agent states that X_t^c received today is as good as m received at t under the condition s^t , we obtain the equality $u'(b^*) \cdot X_t^c = D(t) \cdot p(s^t) \cdot u'(b^*) \cdot m$, and thus,

$$\frac{X_t^c}{m} = D(t) \cdot p(s^t). \quad (2)$$

Since $\frac{X_t^c}{m}$ is observable, we need only elicit $p(s^t)$ to compute the true discount function D . A means of deriving such beliefs is to assume the existence of a commodity h whose utility $v(h)$ is independent of base consumption b_t and additively separable from the utility for money. A possible example of such a commodity could be an anonymous charitable contribution to the charity of the agent’s choice. Then a charity payment of h at time t under the condition s^t is evaluated by¹²

$$D(t) \cdot p(s^t) \cdot v(h).$$

If the subjects exhibits indifference between such a charity payment h at time t under the condition s^t and an unconditional charity payment h made at the same time t but with objective probability α^t , then it is clear from the implied equality $D(t) \cdot p(s^t) \cdot v(h) = D(t) \cdot \alpha^t \cdot v(h)$ that

$$\alpha^t = p(s^t).$$

¹²There is no reason that $v(h)$ must be discounted by D , but this is wlog for our purpose.

Now the *compensated discount function* can be defined by

$$D_c(t) = \frac{X_t^c}{m} \cdot \frac{1}{\alpha^t}, \quad (3)$$

and indeed $D_c(t) = D(t)$.

In what follows we describe the model more precisely.

2.2 Primitives

Time is discrete and with finite horizon, $\mathcal{T} = \{0, 1, \dots, T\}$. The set of possible (inflation-adjusted) future base-consumption levels at time $t > 0$ is given by the finite set $B_t \subset \mathbb{R}_+$ with generic element b_t . This corresponds to assuming the the set of possible income levels are bounded and measurable in the unit of monetary exchange. Because base-consumption will be uncertain, the set B_t is the period t state space. Period 0 base-consumption is given by b^* and $b^* \in B_t$ for all $t \leq T$, that is, period 0 base consumption is known and is a possible consumption level in the future. The t -horizon state space is $\mathcal{S}(t) = \prod_{i=1}^t B_i$, with generic element $s^t = (b_0, \dots, b_t)$. The full state space is $\mathcal{S} = \cup_{t \in \mathcal{T}} \mathcal{S}(t)$.

The set of (inflation-adjusted) monetary prizes is an interval $\mathcal{M} = [0, M]$. Writing $B_t = \{b_t^0, \dots, b_t^{N_t}\}$ with $b_t^0 \leq \dots \leq b_t^{N_t}$, we let M be larger than the grid size for base consumption, $M > \max\{|b_t^i - b_t^{i+1}| : t \leq T, i < N_t\}$. In what follows, we will require that M be small, thereby also requiring a fine state space.

A state-contingent reward is a function $x : \mathcal{S} \rightarrow \mathcal{M}$ that delivers a prize $x(s^t) \in \mathcal{M}$ at date t conditional on the realization of $s^t = (b_0, \dots, b_t)$. The set of all state-contingent rewards is denoted X . The primitive of our analysis is a preference \succsim on X .

2.3 Assumptions

Basic assumptions – The subject is assumed to evaluate future consumption according to a discounted utility model where uncertainty is evaluated according to (state-dependent) subjective expected utility theory. Instantaneous utility is given by an unbounded, strictly increasing and differentiable function $u : \mathbb{R}_+ \rightarrow \mathbb{R}$ with a differentiable inverse. The discount function is $D(t) > 0$ and satisfies $D(0) = 1$ but is not necessarily strictly monotone or restricted to take values less than 1. The subject’s prior (over future base consumption) is a probability measure p on $\mathcal{S}(T)$.

Integration assumptions – Assume that the subject integrates state-contingent rewards with her anticipated base consumption and completely consumes any prize in the period and state that it is received. The presumption here is that the rewards in \mathcal{M} are small enough for this to be an acceptable assumption. It follows that the discounted expected utility due to a state-contingent reward x given beliefs p is

$$U(x) = u(b^*) + \sum_{(b_1, \dots, b_T) \in \mathcal{S}(T)} \left[\sum_{t=0}^T D(t) u(b_t + x(b_0, \dots, b_t)) \right] p(b_0, \dots, b_T). \quad (4)$$

Full support assumption – We assume that $p(s^T) > 0$ for all s^T , that is, unconditional beliefs on each B_t have full support, and positive probability is assigned to future base consumption staying the same as current consumption b^* . Given the strict monotonicity of u , this is equivalent to assuming behaviorally that for any s^T , there exists some prize $m > 0$ such that the state contingent reward x_{m,s^T} that pays m at T in state s^T and 0 otherwise satisfies

$$x_{m,s^T} \succ \omega,$$

where ω denotes the state-contingent rewards that yields 0 at all t and s^t .

2.4 Deriving D from \succsim

Let x_{m,s^t} be the reward that yields prize m at time t conditional on constant base consumption $s^t = (b^*, \dots, b^*)$. Denote by $\psi(x_{m,s^t})$ the reward that yields a prize immediately such that $\psi(x_{m,s^t}) \sim x_{m,s^t}$.¹³ Identify the immediate prize with $\psi(x_{m,s^t})$. The representation (4) implies that

$$u(b^* + \psi(x_{m,s^t})) - u(b^*) = D(t)p_t(s^t)[u(b^* + m) - u(b^*)].$$

Note that since u is a strictly increasing diffeomorphism, $\psi(x_{m,s^t})$ is a strictly increasing differentiable function of m that takes the value 0 when $m = 0$. Taking a derivative of the above expression with respect to m yields

$$u'(b^* + \psi(x_{m,s^t})) \frac{\partial \psi(x_{m,s^t})}{\partial m} = D(t)p_t(s^t)u'(b^* + m).$$

Evaluating at $m = 0$ gives $\frac{\partial \psi(x_{m,s^t})}{\partial m} \Big|_{m=0} = D(t)p_t(s^t) \frac{u'(b^*)}{u'(b^*)}$ and so

$$\frac{\partial \psi(x_{m,s^t})}{\partial m} \Big|_{m=0} = D(t)p_t(s^t).$$

In practice we can rely on an approximation via the observation that for small m ,

$$\frac{\partial \psi(x_{m',s^t})}{\partial m} \Big|_{m'=0} \approx \frac{\psi(x_{m,s^t}) - \psi(x_{0,s^t})}{m - 0} = \frac{\psi(x_{m,s^t})}{m}.$$

Hence

$$\psi(x_{m,s^t}) \approx D(t)p_t(s^t)m$$

Note that $\psi(x_{m,s^t})$ and m are observable, so if we can identify $p_t(s^t)$ (for instance, as described earlier) then we can find $D(t)$.

2.5 Uniqueness of p and D

The primitive preference \succsim on X has the representation (4) where the utility of a reward $x(s^t)$ received at time t conditional on s^t is *state-dependent*.

¹³The existence of such a reward is implied by the unboundedness and continuity of u . Its uniqueness is implied by the strict monotonicity of u .

That is, a dollar received in a given period depends on base consumption b in that period, so that the instantaneous utility in that period is $u(b + 1)$, and thus dependent on b . This might lead one to suspect that the representation (4) lacks desirable uniqueness properties, and in particular the key component of interest, the discount function D , may not be unique. This would be problematic since it would imply that D is not pinned down by preferences \succsim , and in particular, there is no meaningful sense in which it can be extracted from \succsim in any experiment.¹⁴ Therefore, we must establish that any discounted expected utility representation for \succsim must have a unique D . This is the content of the proposition below, the proof of which is relegated to the Appendix.

Proposition 1 *The prior p and the discount function D are uniquely determined by \succsim .*

3 Experimental Design

The experiment used a questionnaire with three sections, all three translated fully in the online appendix. In the first one we asked subjects to choose between rewards received at different times. Some of the future rewards involved payments that were conditional, i.e. only received if the subject’s real income remains “approximately constant” (defined shortly) up to the time of

¹⁴For instance, in state-dependent subjective expected utility, a function f that takes states into prizes is evaluated by $\sum_s u(f(s), s)p(s)$. In this representation, the prior is not unique and thus has no behavioral meaning. We could take any $\alpha_s > 0$ for each s , take a monotone transformation of $u(\cdot, s)$ given by $v(\cdot, s) = \frac{1}{\alpha_s}u(\cdot, s)$ for all s , and adopt a different prior given by $q(s) = \frac{\alpha_s}{\sum_s \alpha_s p(s)}p(s)$ for all s . Then it is easy to see that the utility function $f \mapsto \sum_s v(f(s), s)q(s)$ represents precisely the same preference as before. In contrast, when u is state-independent (as in Savage’s subjective expected utility theory), every subjective expected utility representation for the preference must share the same prior p , and thus p is uniquely pinned down by preferences. It can be elicited by asking the agent to choose between bets.

payment. The second part provided subjects binary choice questions involving a payment to the subject’s charity of choice either with an exogenously given probability or under the condition the subject’s real income remains “approximately constant”. Finally, the third part featured an unincentivized questionnaire eliciting demographic and financial information about subjects, as well as subjective beliefs about future income. All in all, this amounted to six sets of decisions that all included a range of binary choices before they turned to a survey – the small number of questions helps keeps the cognitive burden low on the subjects. During instructions at the beginning of the experiment, we explained to the subjects how a dice roll would determine which one of those six choice categories will be used to generate payoffs, and for that choice one random line from the range of binary choices, according to their expressed preferences.¹⁵

The first section started with questions involving only unconditional payments. A generic question asks for the subject’s preference between¹⁶

(i) a “later” payment of 20,000 Icelandic Kronur (ISK), paid to the subject at t years plus 1 week later,

(ii) a “present” payment in the amount of x paid 1 week later,

where $t=1$ or 2 years, and x ranges from ISK 200 to ISK 22,000 in steps of ISK 200. This series of binary choice questions is presented as a standard multiple price list, that is, in the form of a table.¹⁷ Although subjects could indicate their preference in each cell of the table, for their convenience they were allowed a shortcut where they could indicate two consecutive cells on the

¹⁵It is a dominant strategy to make the binary choices truthfully, with the caveat that if a subject is exactly indifferent between the two options then the choice can be either of them.

¹⁶In 2010, 20,000 ISK was worth approximately \$165.

¹⁷This is equivalent to a Becker-DeGroot-Marschak (BDM) procedure. Like most of the related literature, we opted for the list of binary questions because we believe the original BDM procedure (in particular, understanding why truth-telling is weakly dominant) to be cognitively more demanding for the subjects. This is also in accordance with our experience from several pilots conducted before the experiment.

table where preferences switch from favoring the present reward to preferring the later reward. Consequently, the subject’s indifference point was captured within an ISK 200 interval.

In order to elicit compensated discount functions (and specifically to bring equation (2) into play), we next asked analogous binary choice questions where some payments are paid on the condition that the subject’s income remains “approximately constant”. We formally defined this condition to consist of two requirements:

(a) the price-indexed disposable annual income of the subject during the year following the experiment (in case of a 2-year delayed reward, in both years) is within 4% of the annual income in the 12 months preceding the experiment;

(b) the price-indexed disposable monthly income of the subject in the last month before payment occurs is within 4% of the monthly income in the month preceding the experiment. In the case of payments two years from the experiment, this has to hold true both one year after the experiment and two years after the experiment.

The idea behind part (a) is that the general income level of the subject remains the same, relative to the time of the experiment, while the motivation for part (b) is to make sure that the subject’s overall financial situation is similar to the time of the experiment. We refer to both conditions holding together as the subject’s income situation remaining constant.

Given this definition, the next set of binary choice questions asked subjects to indicate their preference between:

(i’) a “conditional later” payment of 20,000 Icelandic Kronur (ISK) paid to the subject, after t years plus 1 week, *if* income remains approximately constant.

(ii’) a “present” payment in the amount of x paid 1 week later, where as before, $t=1$ or 2 years, and x ranges from ISK 200 to ISK 22,000 in steps of ISK 200.

Section II obtained the data that allows us to exploit equation (3). The section started with the subject choosing a charity from a list of well known and established charities with different objectives that were briefly described to the subjects. Subjects were told that the forthcoming questions involve rewards in the form of payments to their charity of choice. They were then asked to indicate their preference between:

(i) a “conditional charity payment” of ISK 20,000 to the subject’s charity of choice, after time t years plus 1 week, *if* income remains “approximately constant” ,

(ii) a “random charity payment” of ISK 20,000 to the subject’s charity of choice, after time t years plus 1 week, with exogenous probability p , where $t=1$ or 2 years, and p ranges from 0.01 to 1 in steps of 0.01. This series of questions was presented in the form of a table as in Section I.

In Section III, subjects were asked to fill out a survey, which asked them, among other things, for their bank information in order to transfer their payments. Using the one week delay in payments, as opposed to exactly at the time of the experiment, as well as one and two years later, allowed us to use the exact same procedures and payment methods, regardless of whether the rewards were delayed or not. This section also included questions on the subjects’ social and economic background, as well as expectations on their future economic situation. It also included a direct but unincentivized question on how likely they think their yearly income remains approximately the same one and two years following the experiment. We also elicited subjects’ probability assessments of entering a new job by one and two year’s time after the experiment, and their probability assessments on losing their current jobs.

4 Experimental Procedures and Background

4.1 Experimental Sessions

The experiment took place on June 9th and 10th of 2010. Recruitment was conducted by phone from a random sample of Icelanders between the ages of 20 and 45 living in western or central Reykjavik, specifically post codes 101, 105 and 107. The sample was collected from the census by Skýrr, an Icelandic IT company and frequent government contractor. Subsequently, the subjects' phone numbers were collected manually through ja.is, the online Icelandic Telephone Directory. The experiment was conducted in groups simultaneously in one location, specifically a lecture hall at the University of Iceland. The hall had a podium and an overhead projector used in the presentation of instructions, which was carried out by one of the researchers (photos of the location are available upon request). Before starting, the subjects were asked to read and sign a consent form. They were also asked not to talk to each other and informed that if they had questions they should raise their hand, rather than speak up, and they would be assisted individually by a researcher or an assistant. Each session consisted of approximately 15 subjects and took a little bit over one hour. Outside the classroom we set up four dice rolling stations at which assistants reported the randomized outcome of the subjects' dice roll and computerized randomization process.

Before the experimental session in 2010, we conducted several small (10-25 subjects) informal pilots in 2007-2009. These pilots featured similar questions as the experiment, but subjects only received a fixed compensation for participating, independent of their answers. These pilots were mainly used to fine-tune how to effectively explain the questions in the experiment.¹⁸

¹⁸We also conducted two small post-experiment pilots, one in 2011 and one in 2012, with slightly altered questionnaires and instructions as the ones used in the experimental session in 2010, to investigate whether these design changes lead subjects to a better understanding of the questions involving conditional rewards. As we did not find any evidence for this, these post-experimental pilots did not lead to a subsequent incentivized

4.2 Payment Process

As online banking is widespread in Iceland, subjects were paid by bank transfer. In all instances, both for pilots and the experiment itself, subjects received their payments. This happened in the vast majority of instances at the scheduled time. In a few instances with illegible account numbers, payments made with a few days' delay after quick follow-up e-mails or phone calls. Payments were initiated by the finance division at the University of Iceland. In 2010, 48 subjects in the experiment received payments, on average ISK 14,823. In 2011, 17 payments were made, in the amount of ISK 20,673. 12 of those went directly to participants, but five went to charities of their choice. In 2012, 13 individuals received a payment of ISK 21,891 each¹⁹.

4.3 Income Verification

We chose Iceland for the experiment because of the prompt availability of comprehensive income-tax information due to a pay-as-you-earn system where the income tax is continuously withheld at source. That is, the lion's share of income-tax revenue in Iceland is collected monthly, and the Directorate of Internal Revenue (DIR) receives fairly accurate accounts of each individual's income in a timely fashion. For this reason, we signed a contract with the DIR on February 12th 2010, in accordance with the Icelandic Data Protection Act 77/2000 and a notification to the Icelandic Data Protection Authority (S4052). According to this DIR contract and the subjects' informed consent, we did not obtain direct information on subjects' incomes. Instead, income *changes* were calculated by the DIR staff, and they sent us the percentage changes in subjects' monthly as well as yearly incomes, for the specific months and twelve month periods. This was done using the latest

experimental session.

¹⁹These delayed payments were inflation-adjusted equivalents of ISK 20,000 at the time of the experimnt.

information available in the DIR systems, which is generally fairly complete by the 18th of the following month, with only minor adjustments after that.

According to our contract, DIR calculated and delivered the income changes at three points in time. During the week after the experiment took place, as well as one and two years later.²⁰ Income changes were adjusted for changes in the Consumer Price Index (CPI) which is published by Statistics Iceland by the second to last day of the reference month.

4.4 Economic Situation Around the Time of the Experiment

The seemingly flourishing economy of Iceland suffered a major meltdown less than two years before our main experimental session, when the country's three largest banks collapsed and were nationalized. In a widely-viewed televised address, Prime Minister Geir Haarde announced to the country: "(T)here is a very real danger, fellow citizens, that the Icelandic economy, in the worst case, could be sucked with the banks into the whirlpool and the result could be national bankruptcy" (Prime Minister's Office, 2008). Although a sovereign default did not follow, this is indicative of the volatility and uncertainty of the situation. During the following months, hundreds of firms in the country declared bankruptcy. The announcement of the crisis triggered international consequences, including a decision by the United Kingdom to freeze the assets of one of the three large banks (Landsbanki), emergency funding from the International Monetary Fund, protests and a subsequent fall of the government in February 2009.

This was a dramatic macroeconomic shock that affected the entire population of this small open economy with its own currency and for which exchange rates and prices changed suddenly and dramatically. Although the experiment and the post-experimental payments took place after the collapse, and

²⁰We alerted DIR of each deadline with two week's advance notice, and delivery and payments went through with no delay.

our design in principle remains valid no matter what expectations subjects have regarding future income, the crisis should be noted as many subjects may have felt considerable uncertainty as to how the economy would adjust in the coming years. Iceland is one of the world’s smallest currency areas, making the Icelandic krona very vulnerable, which affects the price level and real wages. From the time of the experiment and to the payment dates one and two years later, the price level rose by 4.2% and 9.9% respectively and the real wage rate increased by 2.7% and 4.1% respectively. All in all, it can thus be said that most subjects in our sample faced considerable uncertainty at the time of the experiment regarding their future income.

Our means of deriving compensated discount functions is immune to the economic turbulence as long as subjects use discounted expected utility with respect to some beliefs about the future. To the extent that the ambiguity about the future led subjects to assess the future uncertainty and beliefs in an inconsistent way, it would be revealed to us through unreasonable compensated discount functions.

5 Computation of Discount Factors and Statistical Analyses

Recall our notation X_t^u and X_t^c in Section 2 for the present payments that are indifferent to a later conditional and unconditional payment at t respectively. For simplicity, we estimate these as the lowest payment that the subject indicates as superior to the later payment. For the questions involving only unconditional rewards, if the later payment was paid t years after the experiment, this gives us an uncompensated discount function $D_u(t)$ as defined in equation (1). It will be convenient to define uncompensated discount factors between t and $t + 1$:

$$D_u(t, t + 1) = \frac{D_u(t + 1)}{D_u(t)}.$$

Note that $D_u(0) = 1$ so $D_u(0, 1) = D_u(1)$. When interpreting these parameters in the standard $\beta - \delta$ framework, we set $\delta_u = D_u(1, 2)$, and

$$\beta_u = \frac{(D_u(1))^2}{D_u(2)} = \frac{D_u(1)}{D_u(1, 2)}.$$

To compute compensated discount functions $D_c(t)$, defined by equation (3), we first define p_t as the probability that a subject expects her income t years after the experiment to be approximately the same as at the time of the experiment (as defined in Section 2 and in the experimental instructions). Specifically, this is coded as the lowest reported probability with which a probabilistic contribution to her charity of choice t years after the experiment is indicated as superior to a contribution of the same amount at the same time to the same charity, conditional on her income staying approximately the same. The probabilities computed this way, from incentivized experimental questions, strongly correlate with subjects' reported probabilities of stable income in Section III of the experimental questionnaire, although the latter, unincentivized measures are noisier. Linear regressions of the unincentivized responses on the incentivized ones yield slopes 0.602 and 0.442 with an R^2 of 0.3368 and 0.2154, for the two time horizons respectively.

Given p_t , and X_t^c , which denotes the amount of present payment that makes the subject indifferent to a conditional payment of ISK 20,000 t years after the experiment, the estimated compensated discount function is $D_c(t) = \frac{X_t^c}{p_t \cdot 20,000}$. As with uncompensated discount functions, we define the discount factor $D_c(t, t + 1) = \frac{D_c(t+1)}{D_c(t)}$, $\delta_c = D_c(1, 2)$, and $\beta_c = \frac{(D_c(1))^2}{D_c(2)}$.

For the possible reasons outlined in the introduction – namely the possibility of model misspecification and the possibility of imperfect responses by subjects to the more involved conditional questions – the compensated discount functions computed for a subset of our subjects were not plausible. For example, while it is plausible that a subject may exhibit negative discount rates (that is, $D_c(t) > 1$), it is implausible that discount functions may be steeply upward sloped, as we observed in some cases. Indeed, several

instances of implausible discount functions were typically the result of hard-to-rationalize choices such as preferring a reward m at time t with probability $\alpha < 1$ to a reward m at t for sure.

In light of this, we proceed by following two separate strategies to analyze our predictions for uncompensated and compensated answers. First, for the simpler, uncompensated choices, we report discount functions by strata of income changes, with different predictions under our assumptions that baseline income levels matter for time preferences over monetary rewards. Besides the full sample, those subsamples are (a) people who experienced relatively stable income for two years after the experiment and would thus be expected to show less present bias under the proposed theory, (b) those who have stable income for two years before and after the experiment and are thus assumed to be individuals with even greater stability of income and thus even less confound in the conventional measures of present bias, (c) those who, at the time of the experiment, assess their probability of having a new job to be small and should thus show a smaller present bias than the full sample, and finally (d) those individuals whose realized income rises in the first year after the experiment took place but plateaus after that, who should, according to theory show greater present bias than the full sample or any of the subsamples described above.

Second, for the compensated questions, we identify groups with plausible conditional discount functions, and repeat our analysis restricted to these subsamples. Not to force any arbitrary judgment (and thus results) on the data, we employ the tools of statistical cluster analysis to identify (latent) classes of subjects in the data and focus on groups with a “reasonable” range of discount factors. We choose to define clustering in terms of the discount factors $D_c(0, 1)$ and $D_c(1, 2)$ (the latter only implied by $D_c(1)$ and $D_c(2)$).

For the purposes of clustering, we assume jointly normally distributed discount factors, within an unobserved k class of subjects. Allowing for noise

(outliers), the likelihood function that is numerically maximized is

$$\prod_{i=1}^n \left[\frac{\tau_0}{V} + \sum_{k=1}^G \tau_k \phi_k(\mathbf{x}_i | \theta_k) \right], \quad (5)$$

where V denotes the hypervolume of the data region, data can come from different distributions or simply be noise, which have respective probabilities τ_k (and thus of course $\tau_k \geq 0$ and $\sum_k \tau_k = 1$). The algorithm initializes with noise estimates coming from a nearest-neighbor method and hierarchical clustering applied to the rest of the data (with a simple maximization using the EM algorithm), and the EM algorithm alternating Bayesian updating conditional on the parameter estimates (Expectation step) and maximizing in the parameters conditional on the classification probabilities (Maximization step).

In equation 5 the likelihood contribution for an observation comes from assuming that the densities of the discounts factors follow a bivariate normal distribution. This procedure is conditional on the number of clusters, G . This we let be chosen to maximize the Bayesian Information Criterion (BIC),

$$BIC = 2 \cdot \log L(\mathbf{x}, \theta^*) - (\#\text{parameters}) \log n. \quad (6)$$

Formal hypotheses We conduct a family of tests over two important variants of two measures, using a parametric and a nonparametric test, both two-sided and also one-sided about the economically interesting differences.

Our first set of tests compare first-year and second-year discount factors. The tests are done for both compensated and uncompensated, so we drop the superscript on D below to simplify exposition.

- $H_0 : D(1, 2) = D(0, 1)$, corresponding to subjects being exponential discounters.
- $H_1 : D(1, 2) > D(0, 1)$, corresponding to subjects exhibiting present bias in their time preferences.

We also report test results of no present bias without taking the stance on which way a deviation could occur. This amounts to a two-sided test against the alternative hypothesis of inequality.

Our preferred parametric test is a paired t -test, testing the null hypothesis that $H_0 : E[D(1, 2)] = E[D(0, 1)]$ against the alternative $H_1 : E[D(1, 2)] > E[D(0, 1)]$ or $H_1 : E[D(1, 2)] \neq E[D(0, 1)]$. The test makes the usual assumption that the discount factors are normally distributed in small samples, or the samples are large enough that the asymptotic approximation is good enough. The nonparametric test of our choice is a sign test, testing the medians without assuming the two variables have the same distribution: comparing the null hypothesis that $H_0 : med[D(1, 2)] = med[D(0, 1)]$ against the alternative $H_1 : med[D(1, 2)] > med[D(0, 1)]$ or $H_1 : med[D(1, 2)] \neq med[D(0, 1)]$.

We conduct the tests both for the uncompensated and the compensated discount factors.

Finally, we repeat the exact same test procedures for another parametrization of present bias, namely the transformation of the discount factors into the common $\beta = \frac{D(0,1)}{D(1,2)}$ parameter being larger or less than one.²¹ Thus for the t -tests the null is $H_0 : E[\beta] = 1$ (no present bias on average) against the alternatives of $H_1 : E[\beta] < 1$ or $H_1 : E[\beta] \neq 1$, while for the sign tests, the analogues with medians. Again, we conduct the tests separately for uncompensated or compensated measures.

²¹Of course, for the parametric test the transformation can matter because of the approximation being better or worse in our finite samples. The nonparametric test is indifferent to such a transformation.

6 Experimental Results

6.1 Uncompensated Discount Factors

The bottom pane of Table 1 shows the difference of first and second year uncompensated discount factors ($D_u(1, 2) - D_u(0, 1)$), as well as the implied β and δ parameters, with the strata by income stability around the time of the experiment and times of payment in separate columns. There are 31 subjects in column 2, whose realized annual real income in both years after the experiment stayed within 10% of real income at the time of the experiment. Not surprisingly, these are slightly older subjects, more likely to be employed, but also better educated and with more females than in full sample (The top pane of Table 1 shows summary statistics using our survey in Section III). Reassuringly, these subjects were much less likely to expect job changes after the experiment.²² Column 3 focuses on subjects whose real income was stable both before and after the experiment. Because most of our subjects experienced turbulent income in the years before the experiment, due to the financial crisis and the subsequent recession, we define this category as the subjects whose real income in both years before the experiment was within 20% of real income at the time of the experiment, and whose real income in both of the two years after the experiment was within 10% of real income at the time of the experiment. Even this more permissive criterion for the two years before the experiment results in only a small number of subjects (10, with 9 females) being in this category. They are also better educated and slightly older.

²²For those in Column 2 whose realized income stayed stable for two years after the experiment, the assessment (at the time of the experiment) of being in a new job one year (two years) after the experiment is 20.0 (25.6) percentage points less likely than those not having such stable income, and this difference has a p -value of 0.003 (0.00). For those in column 3 who we categorized as having stable income both before and after the experiment, these differences from the rest of the sample are 26.8 and 27.1 in the same direction (with both p -values of 0.00).

We assume that those whose realized income was stable for two years after the experiment expected more stable income in these years than the rest of the subjects. This is consistent with subjects' reported beliefs in Section III of the questionnaire, and the highly correlated incentivized responses in Section II. Yet if we define stable income subjects directly based on their reported expectations in the questionnaire, instead of their realized income that we directly observe, we get results similar to what follows, but noisier (Column 4 focuses on the 32 people who estimated the risk of have a new job in two years to be less than 20%). This suggests that at least some subjects have trouble reporting their expectations in terms of probabilities.

As shown in the bottom pane of Table 1, the difference between the second and first year discount factors is large and highly statistically significant for the entire sample of 115 (we lose one observation with an invalid second-period discount factor). The overall sample shows considerable present-bias (average β is 0.895), and a reasonable amount of impatience (on average δ being 0.922). In contrast to this, for those 31 people whose incomes were stable for two years after the experiment, the difference in discount factors is statistically insignificant, with a point estimate close to 0, and the implied β is only insignificantly below 1, with a point estimate of 0.972. On the other hand, the estimated δ (0.924) is essentially the same as for the whole subject pool. We see the same pattern for those subjects whose income was stable both before and after the experiment.

We can contrast these results with the 11 subjects who experienced a relatively large real income rise (more than 10%) in the first year after the experiment, but then saw their incomes stabilized (remained within 10% of the income in the first year after the experiment). If these subjects foresaw this income path,²³ their expected marginal utilities for monetary rewards in both of the two years following the experiment are lower than at the time of

²³This is reasonable in many instances, e.g. for those finishing school or for some other reason expecting to get into a higher-paying job.

the experiment. This should imply more present bias than for the rest of the subject pool (without compensating for the confound of the income process, of course). We find some evidence for this, as the estimated $D_u(1, 2) - D_u(0, 1)$ difference rises to 0.223, and the estimated β decreases to 0.792 for this subsample, although we have no power in this small sample to establish these changes as statistically significant.

6.2 Compensated Discount Factors

6.2.1 Clustering outcomes

The noise and clusters resulting from the procedure in Section 5 are summarized in the top pane of Table 2. The best-fitting distribution is one with no correlations between the two discount rates in any cluster, and three clusters with different means and variances for the normal distributions, and some outliers (noise).²⁴ Descriptive statistics for clusters are shown in the middle pane of Table 2. The first two clusters, containing 79 of the 116 subjects, have compensated discount factors mostly in a plausible range. Since Cluster 1 still contains some individuals with high compensated discount factors for some period of time, we also focus on cluster 2 separately, although this group only contains 16% of our sample. The “reasonable” clusters are slightly better educated, on average (see Table 2), which is consistent with the hypothesis that these subjects understood the questions involving conditional rewards better, but the overlap between all clusters is apparent for any covariate.

²⁴The conclusions are robust to different initializations of the clustering. Though the number of clusters might be higher with a different initialization (e.g. initializing the expected outliers differently), the clusters with “reasonable” discount factors largely overlap, as do the vector of noise-indicators.

6.2.2 Main results

Comparisons of uncompensated and compensated discount factors are reported in the bottom pane of Table 2, with standard errors in parentheses. The p-values of our tests outlined above are reported in Table 3. As with the average discount factors before, some statistics for the whole dataset are quantitatively implausible, although they imply the same qualitative conclusions as we obtain when restricting attention to subjects in the two reasonable clusters. We find the same clear patterns no matter whether we include all subjects, only those in Clusters 1 and 2, or only those in Cluster 2. Compensated discount factors, on all horizons are significantly higher than uncompensated ones: $D_c(0, 1) > D_u(0, 1)$ and $D_c(1, 2) > D_u(1, 2)$. This suggests that standard elicitation of time preferences, which do not take into account future income expectations of subjects, on average overestimate the amount of impatience of individuals. One explanation for this is that most subjects expect a rising income path, decreasing the marginal utility of small monetary rewards in the future, relative to the present.

Second, the left pane of Table 3 shows that our tests reject exponential discounting without compensation for income changes in the entire sample, but not for the groups with stable income, restating the conclusions of our previous analysis reported in Table 1. More interestingly, the new results on compensated discount factors in the right pane of Table 3 show no evidence of present bias once we control for the potential confounds of income changes. This suggests that once one compensates for income expectations, there is insufficient evidence that the average subject has a present bias.

7 Conclusion

We find that in a setting where most people reveal present bias, people with stable incomes are standard exponential discounters. Our results show that

subjects' choices over delayed rewards depend on their income expectations, revealing their importance for intertemporal choice experiments. Our experimental design shows how choices on rewards conditional on no income changes, alongside incentivized elicitations of subjective income expectations, can help researchers compensate for this confound. While we had the maintained hypothesis that personal annual income equals personal consumption in the same period, our results cast doubt on the maintained hypothesis of many other studies that mental accounting implies that subjects evaluate monetary rewards without integrating them to other income.

It is important for future research to obtain a richer empirical picture on how time preferences over monetary payments depend on different characteristics of future income expectations such as trends in expected income, trends in volatility of income, or the amount of autocorrelation of income in future periods. Any conclusion on saving behavior or policy would be overly speculative at this point, but for fundamental work in choice theory, this is an important step towards more conclusive lab experiments and empirical studies.

8 References

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

Andreoni, J. and C. Sprenger (2012): “Estimating time preferences from convex budgets,” *American Economic Review* 102, 3333-3356.

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

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

Carvalho, L., S. Meier and S. Wang (2014): “Poverty and economic decision-making: Evidence from changes in financial resources at payday,” *mimeo Columbia University*.

Dean, M. and A. Sautmann (2014): “Credit constraints and the measurement of time preferences,” *mimeo Brown University*.

Eppel, T. (2010): “Rational planners or myopic fools?,” *mimeo ETH Zurich*.

Fernandez-Villaverde, J. and A. Mukherji (2011): “Can we really observe hyperbolic discounting?,” *mimeo University of Pennsylvania*.

Fraley, C. and A. Raftery (2006): “MCLUST Version 3 for R: Normal mixture modeling and model-based clustering,” *Technical Report no. 504, Department of Statistics, University of Washington*.

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

Fudenberg, D. and D. Levine (2006): “A dual-self model of impulse control,” *American Economic Review* 96, 1449-1476.

Gerber, A. and K. Rohde (2010): “Risk and preference reversals in intertemporal choice,” *Journal of Economic Behavior and Organization* 76,

654-668.

Gine, X., J. Goldberg, D. Silverman and D. Yang (2013): “Revising commitments: Field evidence on the adjustment of prior choices,” *mimeo World Bank*.

Gul, F. and W. Pesendorfer (2001): “Temptation and self-control,” *Econometrica* 69, 1403-1435.

Halevy, Y. (2012): “Time consistency: Stationarity and time invariance,” *mimeo University of British Columbia*.

Harrison, G., M. Lau and M. Williams (2002): “Estimating individual discount rates in Denmark: A field experiment,” *American Economic Review* 92, 1606-1617.

Harrison, G., M. Lau and E. Rutstrom (2005): “Dynamic consistency in Denmark: Alongitudinal field experiment,” *mimeo University of Central Florida*.

Krupka, E. and M. Stephens (2013): “The stability of measured time preferences,” *Journal of Economic Behavior and Organization*, 85, 11-19.

Laibson, D. (1997): “Golden eggs and hyperbolic discounting,” *Quarterly Journal of Economics*, 112, 443-477.

McClure S, Ericson K, Laibson D, Loewenstein G, Cohen J. (2007) “Time Discounting for Primary Rewards.” *Journal of Neuroscience*. 27(21):5796–5804.

McLachlan, G. and T. Krishnan (2008): *The EM algorithm and extensions*; Hoboken, New Jersey: Wiley.

McLachlan, G. and D. Peel (2000): *Finite mixture models*; New York: Wiley.

Meier, S. and C. Sprenger (2010): “Stability of time preferences,” *mimeo Columbia University*.

Noor, J. (2009): “Hyperbolic discounting and the standard model: Eliciting discount functions,” *Journal of Economic Theory* 144, 2077-2083.

O’Donoghue, T. and M. Rabin (1999): “Doing it now or later,” *American*

Economic Review 89, 103-124.

Olson, M. and M. Bailey (1981): "Positive time preference," *Journal of Political Economy*, 89, 1-25.

Pender, J. (1996): "Discount rates and credit markets: Theory and evidence from rural India," *Journal of Development Economics* 50, 257-296.

Rubinstein, A. (2002): "Comments on the risk and time preferences in economics," *mimeo Tel Aviv University*.

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

A Appendix A: Proof of Proposition

Proof. Denote by p_t the probability measure over $\mathcal{S}(t)$ induced by p . We first show that p_t is pinned down by \succsim . Write $B_t = \{b^0, \dots, b^N\}$ such that $b^0 \leq \dots \leq b^N$ (the t subscripts are dropped to ease notation). For any $i \leq N$, let x_{m,s^t} be the reward that yields prize m at time t conditional on $s^t = (s_1, \dots, s_{t-1}, b^i)$. Denote by $\psi(x_{m,s^t})$ the reward that yields a prize immediately such that $\psi(x_{m,s^t}) \sim x_{m,s^t}$.²⁵ Identify the immediate prize with $\psi(x_{m,s^t})$. The representation implies that

$$u(b^* + \psi(x_{m,s^t})) - u(b^*) = D(t)p_t(s_1, \dots, s_{t-1}, b^i)[u(b^i + m) - u(b^i)].$$

Note that since u is a strictly increasing diffeomorphism, $\psi(x_{m,s^t})$ is a strictly increasing differentiable function of m that takes the value 0 when $m = 0$. Taking a derivative of the above expression wrt m yields

$$u'(b^* + \psi(x_{m,s^t})) \frac{\partial \psi(x_{m,s^t})}{\partial m} = D(t)p_t(s_1, \dots, s_{t-1}, b^i)u'(b^i + m).$$

Evaluating at $m = 0$ gives

$$\left. \frac{\partial \psi(x_{m,s^t})}{\partial m} \right|_{m=0} = D(t)p_t(s_1, \dots, s_{t-1}, b^i) \frac{u'(b^i)}{u'(b^*)}. \quad (7)$$

Consider the contingent reward $x_{m,s^t}^{b^1-b^0}$ that gives $b^1 - b^0$ at time t unconditionally and in addition gives m at time t conditional on $s^t = (s_1, \dots, s_{t-1}, b^0)$. Denote by $\psi(x_{m,s^t}^{b^1-b^0})$ the contingent prize that pays $b^1 - b^0$ at time t unconditionally and also an immediate prize (which, abusing notation, is also denoted by $\psi(x_{m,s^t}^{b^1-b^0})$) satisfies $\psi(x_{m,s^t}^{b^1-b^0}) \sim x_{m,s^t}^{b^1-b^0}$. Then by the representation,

$$\begin{aligned} u(b^* + \psi(x_{m,s^t}^{b^1-b^0})) - u(b^*) &= D(t)p_t(s_1, \dots, s_{t-1}, b^0)[u(b^0 + b^1 - b^0 + m) - u(b^0 + b^1 - b^0)] \\ &= D(t)p_t(s_1, \dots, s_{t-1}, b^0)[u(b^1 + m) - u(b^1)]. \end{aligned}$$

²⁵The existence of such a reward is implied by the unboundedness and continuity of u . Its uniqueness is implied by the strict monotonicity of u .

Moreover,

$$\frac{\partial \psi(x_{m,s^t}^{b^1-b^0})}{\partial m} \Big|_{m=0} = D(t) p_t(s_1, \dots, s_{t-1}, b^0) \frac{u'(b^1)}{u'(b^*)}.$$

Similarly, for each $0 < i \leq N$,

$$\frac{\partial \psi(x_{m,s^t}^{b^i-b^{i-1}})}{\partial m} \Big|_{m=0} = D(t) p_t(s_1, \dots, s_{t-1}, b^{i-1}) \frac{u'(b^i)}{u'(b^*)}. \quad (8)$$

By assumption, $p_t(s^t) \neq 0$ for all s^t (that is, $\frac{\partial \psi(x_{m,s^t})}{\partial m} \Big|_{m=0} \neq 0$ for all s^t). Using (7) and (8), we get

$$\frac{\frac{\partial \psi(x_{m,s^t}^{b^i-b^{i-1}})}{\partial m} \Big|_{m=0}}{\frac{\partial \psi(x_{m,s^t}^{b^1-b^0})}{\partial m} \Big|_{m=0}} = \frac{p_t(s_1, \dots, s_{t-1}, b^{i-1})}{p_t(s_1, \dots, s_{t-1}, b^i)}$$

for all i . That is, the noted relative probabilities are determined uniquely by preferences. Conclude that $p_t(s^t)$ is uniquely determined for each s^t . In particular, p is uniquely determined.

To see that D is uniquely pinned down by preferences, note that by (7), for $b^i = b^*$ (that is, $s^t = (s_1, \dots, s_{t-1}, b^0)$),

$$\frac{\partial \psi(x_{m,s^t})}{\partial m} \Big|_{m=0} = D(t) p_t(s^t). \quad (9)$$

Thus, the uniqueness of $p_t(s^t)$ implies that of $D(t)$. ■

Table 1: Description of the Sample

	Full sample	Stable income after 2010 ^y	Stable income before and after 2010 ^{yy}	New job less likely than 20% ^{yyy}	Income rises, then plateaus ^{yyyy}
Summary Statistics (with SDs)					
Female	0.59 (0.49)	0.68 (0.48)	0.90 (0.32)	0.59 (0.50)	0.64 (0.50)
Age	33.19 (6.90)	36.68 (6.37)	37.20 (5.35)	35.66 (6.99)	32.18 (7.36)
Higher education	0.43 (0.50)	0.55 (0.51)	0.70 (0.48)	0.44 (0.50)	0.45 (0.52)
Smoker	0.20 (0.40)	0.23 (0.43)	0.20 (0.42)	0.13 (0.34)	0.27 (0.47)
BMI	25.43 (4.52)	24.55 (3.48)	23.93 (2.77)	25.34 (4.27)	25.35 (4.27)
Employed	0.88 (0.33)	0.94 (0.25)	0.90 (0.32)	0.88 (0.34)	1.00 (0.00)
Single	0.60 (0.49)	0.74 (0.44)	0.80 (0.42)	0.72 (0.46)	0.45 (0.52)
Number of children	1.22 (1.36)	1.35 (1.36)	1.30 (1.42)	1.78 (1.48)	1.45 (1.57)
Expect job change in t=1	35.44 (35.46)	20.73 (28.74)	11.00 (15.78)	2.78 (3.96)	20.09 (24.72)
Exact job change in t=2	47.76 (36.64)	28.93 (30.55)	23.00 (26.37)	4.91 (5.40)	42.73 (36.36)
Time Preference Variables (with SEs)					
$D_u(1,2)-D_u(0,1)^*$	0.13 (0.03)	0.04 (0.04)	0.03 (0.06)	0.08 (0.04)	0.22 (0.10)
β_u^{**}	0.89 (0.03)	0.97 (0.04)	0.98 (0.06)	0.94 (0.04)	0.79 (0.10)
δ_u^{***}	0.92 (0.02)	0.92 (0.02)	0.92 (0.03)	0.90 (0.03)	0.89 (0.06)
N	116	31	10	32	11

Note: The table reports means and standard deviations (in parentheses) for basic descriptive variables for various subsamples in the top pane, means and standard errors (in parentheses) for time preference variables in the bottom pane. * Reports the average difference between each subject's revealed uncompensated discount factor over the second year and the one over the first year. ** Reports the average ratio of discount factors over the second year and over the first year. *** Reports the average ratio of the two-year discount factor and the one-year discount factor. ^y Only including those whose realized annual income remained within 10% of their 2010 income both in 2011 and 2012. ^{yy} Only including those whose annual incomes in 2008 and 2009 as well. ^{yyy} Includes those whose self-report the probability of getting a new job in either year after the experiment to be less than 20%. ^{yyyy} Includes those whose realized annual income rose more than 10% in 2011 but less than 10% on top of that in 2012.

Table 2: Parameters by Clusters

	Full sample	Cluster 1	Cluster 2	Clusters 1&2	Cluster 3	Noise
Estimated Parameters of Clustering						
Probability*		0.497	0.159		0.198	0.147
$E[D_c(0,1)]^{**}$		1.374	1.008		1.692	
$V[D_c(0,1)]^{***}$		0.670	0.001		2.702	
$E[D_c(1,2)]^{**}$		1.150	0.971		4.107	
$V[D_c(1,2)]^{***}$		0.183	0.005		9.123	
Summary Means (SDs)						
Female	0.59 (0.49)	0.63 (0.49)	0.53 (0.51)	0.61 (0.49)	0.52 (0.51)	0.56 (0.51)
Age	33.19 (6.90)	33.48 (7.29)	31.84 (6.27)	33.09 (7.05)	32.62 (7.02)	34.44 (6.23)
Higher education	0.43 (0.50)	0.47 (0.50)	0.53 (0.51)	0.48 (0.50)	0.43 (0.51)	0.19 (0.40)
Smoker	0.20 (0.40)	0.23 (0.43)	0.16 (0.37)	0.22 (0.41)	0.14 (0.36)	0.19 (0.40)
BMI	25.43 (4.52)	25.22 (4.55)	25.32 (5.85)	25.24 (4.87)	25.20 (2.98)	26.67 (4.59)
Employed	0.88 (0.33)	0.88 (0.33)	0.84 (0.37)	0.87 (0.34)	0.81 (0.40)	1.00 (0.00)
Single	0.60 (0.49)	0.62 (0.49)	0.74 (0.45)	0.65 (0.48)	0.52 (0.51)	0.50 (0.52)
Number of children	1.22 (1.36)	1.25 (1.36)	1.11 (1.15)	1.22 (1.31)	1.00 (1.34)	1.50 (1.63)
Expect job change in t=1	35.44 (35.46)	27.52 (30.31)	33.68 (38.87)	29.04 (32.47)	50.62 (38.19)	46.31 (39.10)
Expect job change in t=2	47.76 (36.64)	39.76 (34.63)	43.42 (44.07)	40.66 (36.90)	65.24 (30.39)	59.00 (34.09)
Sample Estimates (SEs)						
$D_c(0,1)$	5.963 (1.666)	1.399 (0.108)	1.005 (0.007)	1.305 (0.084)	1.767 (0.408)	34.471 (9.539)
$D_c(1,2)$	2.832 (0.530)	1.145 (0.056)	0.975 (0.016)	1.104 (0.043)	4.510 (0.682)	9.161 (3.296)
$D_c(0,1)-D_u(0,1)$	5.171 (1.664)	0.578 (0.103)	0.226 (0.046)	0.493 (0.081)	1.037 (0.396)	33.692 (9.519)
$D_c(1,2)-D_u(1,2)$	1.926 (0.536)	0.237 (0.056)	0.092 (0.037)	0.202 (0.044)	3.541 (0.672)	8.212 (3.323)
$D_c(1,2)-D_c(0,1)$	-3.131 (1.803)	-0.254 (0.123)	-0.030 (0.018)	-0.200 (0.094)	2.743 (0.992)	-25.310 (11.838)
$D_u(1,2)-D_u(0,1)$	0.130 (0.027)	0.088 (0.031)	0.104 (0.053)	0.092 (0.027)	0.240 (0.084)	0.171 (0.080)
β_u	0.895 (0.026)	0.929 (0.031)	0.899 (0.056)	0.922 (0.027)	0.791 (0.079)	0.898 (0.082)
β_c	16.075 (8.997)	1.487 (0.175)	1.037 (0.020)	1.379 (0.135)	1.634 (0.714)	107.595 (62.053)
N	116	60	19	79	21	16

Note: The table collects the parameters corresponding to the clusters as well as means and standard errors (in parentheses) for key constructs. *Estimated probabilities of the cluster in the mixture model. **Estimated mean of the compensated discount factors. *** Estimated variance for the discount factors.

Table 3: P-values from Hypothesis Tests

Uncompensated Discount Factors					
Null hypothesis	Test	Full sample	Stable income after 2010 ^y	Stable income before and after 2010 ^{yy}	Income rises, then plateaus ^{yyy}
$D_u(1,2) \leq D_u(0,1)$	Student's t	0.000	0.161	0.318	0.023
	sign test	0.032	0.828	0.945	0.113
$\beta_u \geq 1$	Student's t	0.000	0.242	0.368	0.034
	sign test	0.032	0.828	0.945	0.113
$D_u(1,2) - D_u(0,1) = 0$	Student's t	0.000	0.322	0.637	0.046
	sign test	0.064	0.572	0.344	0.227
$\beta_u = 1$	Student's t	0.000	0.483	0.736	0.067
	sign test	0.064	0.572	0.344	0.227
N		115	31	10	11
Compensated Discount Factors					
		Full sample	Cluster 1	Cluster 2	Clusters 1&2
$D_u(1,2) \leq D_u(0,1)$	Student's t	0.957	0.979	0.948	0.982
	sign test	0.719	0.817	0.945	0.906
$\beta_u \geq 1$	Student's t	0.952	0.996	0.958	0.997
	sign test	0.719	0.817	0.945	0.906
$D_u(1,2) - D_u(0,1) = 0$	Student's t	0.085	0.043	0.105	0.036
	sign test	0.699	0.519	0.344	0.282
$\beta_u = 1$	Student's t	0.097	0.007	0.083	0.006
	sign test	0.699	0.519	0.344	0.282
N		115	60	19	79

Note: The table collects p-values for tests with the null hypothesis corresponding to present bias. Low p-values thus indicate hyperbolicity or simply "non-exponentiality". ^y Only including those whose realized annual income remained within 10% of their 2010 income both in 2011 and 2012. ^{yy} Only including those whose annual incomes in 2008 and 2009 as well. ^{yyy} Includes those whose realized annual income rose more than 10% in 2011 but less than 10% on top of that in 2012.