University of Neuchatel

Institute of Economic Research

IRENE Working paper 17-02

Empirical identification of time preferences: Theory and an illustration using convex time budgets

Antoine Bommier

Bruno Lanz

UNIVERSITÉ DE NEUCHÂTEL

Institut de recherches économiques

Empirical identification of time preferences: Theory and an illustration using convex time budgets

Antoine Bommier[*] Bruno Lanz[†]

This version: December 2017

Abstract

We develop a simple theoretical framework that identifies time preferences without relying on a particular utility function. Our empirical strategy requires observations about intertemporal consumption allocation decisions made under varying relative prices, and seeks to approximate the marginal rate of substitution of consumption at different dates along a constant consumption path. Doing so, we emphasize the importance of measuring the curvature of the intertemporal utility function (or willingness to substitute consumption across time). We illustrate our approach with data derived from the convex time budget procedure of Andreoni and Sprenger (AER, 2012).

Keywords: Intertemporal choice; Discounting behavior; Intertemporal substitution; Discounted utility model; Convex budgets

JEL classification: D03, D12, D91, E61

[*]ETH Zürich, Chair for Integrative Risk Management and Economics, Switzerland. E-mail: abommier@ethz.ch.
[†]University of Neuchâtel, Department of Economics and Business, Switzerland. E-mail: bruno.lanz@unine.ch.

1 Introduction

In an intertemporal setting, the discount factor measures the marginal rate of substitution (MRS) at two dates, holding consumption constant (e.g. Epstein, 1987). In principle, time preferences can therefore be empirically identified by observing relative prices at which consumption is constant over time (all other things equal). Because this is typically hard to observe, current estimate of time preferences exploit data on non-constant intertemporal consumption allocations to extrapolate the MRS that would prevail along a constant consumption trajectory. One popular empirical strategy to approach this problem is to specify a parsimonious functional form for the intertemporal utility function, so that estimating the parameters of that function from *local* consumption allocation can be used to approximate *global* properties of intertemporal preferences.

In an important paper, Andersen et al. (2008) show imposing assumptions about the intertemporal utility function can have major implications for estimates of time preferences. Relying on two multiple price list (MPL) tasks, separately eliciting preferences over intertemporal trade-offs and static lotteries, they develop a maximum likelihood procedure to jointly estimate the rate of time preferences and the coefficient of relative risk aversion. Their results suggest that imposing linear (risk-neutral) per-period utility biases estimates of the discount factor downwards. In a similar vein, Andersen et al. (2011) use a third MPL task to control for a measure of intertemporal risk aversion (or correlation aversion), another source of curvature for the intertemporal utility function. In another important contribution, Andreoni and Sprenger (2012) suggest an alternative experimental procedure that allows subjects to allocate a fixed amount of consumption "tokens" across two different dates in response to varying relative prices, thereby convexifying the choice problem. This "convex time budgets" (CTB) procedure allows jointly measuring time preferences and the curvature of the utility function.

While both approaches emphasize the need to control for the curvature of the utility function in order to identify time preferences, they rely on a different concept of curvature. In particular, while Andersen et al. (2008) use risk preferences, Andreoni and Sprenger (2012) employ a measure of preferences for consumption smoothing, or intertemporal elasticity of substitution (IES). But despite this conceptual difference, estimation in both papers impose a specific intertemporal utility function to identify the parameters of interest.

In this paper we take a different route to tackle the same problem. We start from the observation that variations in the MRS are observed when individuals make consumption allocation choices at different relative prices. Using a simple theoretical framework, we first show how the MRS on the relevant point in the consumption space (i.e. when consumption is constant) can be approximated from non-constant consumption tuples. Importantly, our approximation is based on a general measure of the curvature of intertemporal indifference curves, it does not necessitate the use of a specific functional form. In other words, while our approach also emphasizes the role of the curvature of the utility function, we remain agnostic about the parametric form of the intertemporal utility function. It is also worth noting that the setup we consider does not involve choice under uncertainty, so that the curvature of the intertemporal utility function does not quantify risk preferences.[1]

To operationalize our approach, we need to observe how individuals allocate consumption at two different dates under varying relative prices. Therefore, a particularly relevant data generation process is provided by the CTB procedure of Andreoni and Sprenger (2012), which supplies exactly the co-variations we require. Hence in the second part of the paper, we make use the data of the original contribution by Andreoni and Sprenger (2012) data to illustrate our approach. Interestingly, applying our approach to this particular dataset is also relevant, as Echenique et al. (2015) provide evidence that a large proportion of observed choices fail to comply with some basic implication of the standard additive model.[2] This makes our approach particularly attractive as it does not impose specific assumptions about the utility function representing intertemporal preferences.

Our empirical results suggest that substitutability of consumption at two different dates is very high. Given the relatively short horizon considered in the experiment (between 5 and 14

[1] Methods that use choice under uncertainty to estimate the rate of time preferences rely on assumptions regarding risk preferences. However, since the rate of time preference is a property of ordinal preferences, thus directly reflected in how people make choices in absence of uncertainty, these additional assumptions can in principle be avoided.

[2] More specifically, in the data collected by Andreoni and Sprenger (2012), thirty percent of subjects are consistent with an assumption of exponential discounting and less than fifty percent of subjects are consistent with more general time-additive preferences.

weeks separate the first from the second payments), this result seems intuitive. High intertemporal substitutability is also reflective of the fact that many subjects chose corner allocations (where the entire budget is allocated to either the first or the second date). Our estimates of intertemporal substitutability are in fact significantly larger than those reported in Andreoni and Sprenger (2012). An implication of lower utility curvature, which is in line with the argument made in Andersen et al. (2008), is that our estimates of the rate of time discounting are also larger than those reported in Andreoni and Sprenger (2012). Moreover, our approach suggests that intertemporal substitutability increases both with consumption and decreases with the delay between consumption dates, which is inconsistent with the standard additive model with constant relative risk aversion (CRRA) per-period utility.

In sum, by developing a novel approximation approach to the estimation of time preferences, our work makes two contributions to the empirical literature on time preferences. First, we highlight that the rate of time preference can be identified without reference to a specific functional form for the utility function. As initially highlighted by Andersen et al. (2008), adhoc structural assumption may bias estimates of the parameters of interest. Second, we clarify that the ideal data generation process to identify time preference should yield variations in consumption choices that are not "too far" away from a constant consumption trajectory. The fact that the CTB data features many corner outcomes, which lie relatively far away from the point at which time preferences can be measured, implies that the choice of the functional form will matter a lot. This is because global properties of the function have to be inferred from boundary outcomes. Our work therefore suggests that further refinements in procedures to elicit time preferences are warranted.

The remaining of the paper is organized as follows. In Section 2 we present our basic theoretical framework. Section 3 describes our empirical illustration. This includes a brief description of Andreoni and Sprenger (2012), our maximum likelihood estimation strategy, and estimation results. Section 4 discusses our findings and concludes.

2 Theory

In this section we first introduce the notation and two useful definitions to measure time preferences and substitutability between consumption in different periods. We then provide the main result of the paper on how time preferences can be identified from observations on relative prices and consumption choices, without imposing a specific functional form to represent intertemporal preferences.

2.1 Notation and definitions

We consider a consumption allocation problem between two periods $t, t' \in 1, 2$, where consumption of a good (money) is written c_t. Preferences for consumption are represented by a utility function $U(c_1, c_2)$. We denote the partial derivative with respect to argument $t = 1, 2$ by $\partial U(c_1, c_2)/\partial c_t = U_t(c_1, c_2)$, or just U_t whenever it does not create confusion. Similarly, we denote $\partial^2 U(c_1, c_2)/\partial c_t c_{t'} = U_{tt'}(c_1, c_2)$ or $U_{tt'}$. We assume that these partial derivatives exist.

We define the discount rate as a measure of how rapidly the marginal utility of consumption declines with time, controlling for variation in consumption (see Epstein, 1987, for a similar definition). Formally, the discount factor and the pure rate of time preferences can then be defined as follows.

Definition 1. *The discount factor is:*

$$\delta = \left. \frac{\partial U/\partial c_2}{\partial U/\partial c_1} \right|_{c_1=c_2} \tag{1}$$

and the pure rate of time preference associated with the discount factor δ is:

$$\rho = (1/\delta) - 1 \,. \tag{2}$$

This definition of time preferences describes a property of intertermporal preferences rather than a parameter of a particular function representing the associated preference ordering. Intuitively, it is a measure of the slope of indifference curves at points where $c_1 = c_2$, and represents the willingness to trade-off consumption at both dates, as represented by the MRS between time-dated consumption goods.

The second definition provides an absolute measure of the curvature of the intertemporal utility function.

Definition 2. *The (symmetric) coefficient of absolute intertemporal substitutability (CAIS) is:*

$$\gamma(c_1, c_2) \;=\; \frac{1}{-\frac{U_{11}}{U_1} + U_{12}\left(U_1^{-1} + U_2^{-1}\right) - \frac{U_{22}}{U_2}} \tag{3}$$

This quantity summarizes second order properties of the intertemporal utility function and is similar in spirit to the direct elasticity of substitution between consumption in two periods, introduced by McFadden (1963), which is also known as the IES. As for our definition of time preferences, the CAIS is an ordinal property and hence is independent of the choice of a particular utility function to represent the preference relation (see Appendix).

To concretely illustrate these definitions, we consider the standard additive separable specification

$$U(c_1, c_2) = u(c_1) + \beta u(c_2)\,,$$

where $u(\cdot)$ is the per period utility function. First, time preferences are of course given by $\delta = \beta$, which is commonly called the discount factor. Furthermore, by construction we have that $U_{12} = 0$, so that the CAIS reduces to:

$$\gamma(c_1, c_2) \;=\; -\frac{U_{11}}{U_1} - \frac{U_{22}}{U_2}\,.$$

Further assuming that $u(c_t) = \frac{c_t^{1-\eta}}{1-\eta}$, the CAIS in the additive model is proportional to the direct elasticity of substitution evaluated when $\delta = 1$ and $c_1 = c_2$:

$$\gamma = \frac{2}{c_t \eta}\,. \tag{4}$$

Note that given this functional form γ declines in c_t, whereas the IES is a constant and does not vary with consumption.

2.2 Identification of time preferences

Maximizing an intertemporal utility function $U(c_1, c_2)$ subject to relative prices p_1/p_2 (and income m) yields the following condition for an interior solution:

$$\frac{p_1}{p_2} = \frac{U_1(c_1^*, c_2^*)}{U_2(c_1^*, c_2^*)}, \tag{5}$$

where $MRS(c_1^*, c_2^*) = \frac{U_1(c_1^*, c_2^*)}{U_2(c_1^*, c_2^*)}$. Given definition 1, measuring time preferences coincides with estimating the slope of an indifference curve at a point where $c_1^* = c_2^*$. Hence if we could observe the price ratio that would induce consumers to select some point (c^*, c^*) it would directly reveal time preferences.

However in most cases we observe consumptions choices for which $c_1^* \neq c_2^*$, so that the price ratio associated with consumption choices is equal to the MRS away from the 45 degree line. We can, however, employ observations of the MRS at observed points (c_1, c_2) in order to approximate its value at a nearby point where $c_1 = c_2$. One natural candidate to do such an approximation is to use average consumption $c^* = \frac{c_1^* + c_2^*}{2}$, as it is the projection of observed consumption (c_1, c_2) on the 45° line, and is thus the "closest" relevant point from the observed data point. As shown in Figure 1, the distance from (c_1^*, c_2^*) to (c^*, c^*) is then $\varepsilon = \frac{c_1^* - c_2^*}{2}$.

One approach to approximate the MRS on the 45 degree line given observations away from it is to use a Taylor expansion. Formally, for the first period, we have:

$$U_1(c_1^*, c_2^*) \simeq U_1(c^*, c^*) + \varepsilon U_{11}(c^*, c^*) - \varepsilon U_{12}(c^*, c^*) .$$

Similarly, for second period consumption c_2 we have:

$$U_2(c_1^*, c_2^*) \simeq U_2(c^*, c^*) + \varepsilon U_{12}(c^*, c^*) - \varepsilon U_{22}(c^*, c^*) .$$

Forming the ratio of these two expressions, we obtain an approximation of the MRS, evaluated

Figure 1: Intertemporal consumption choice and convex budget

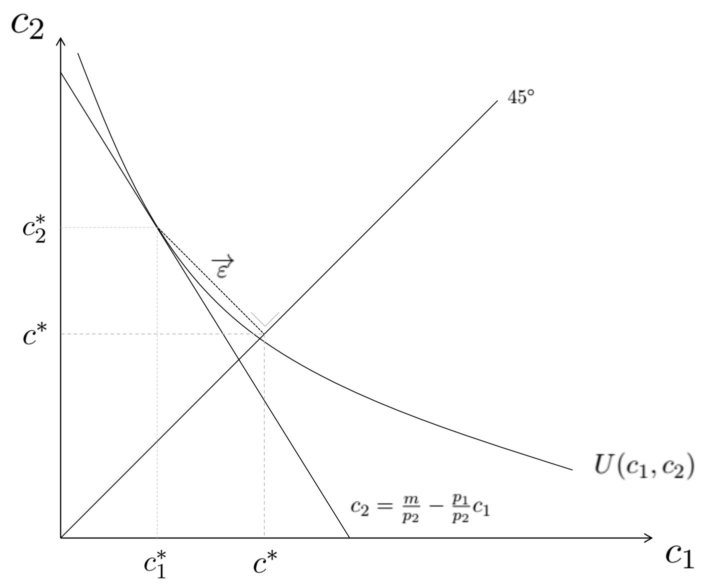

at (c_1, c_2), as a function of c^*. In particular, from Equation (5), we get that:

$$\frac{p_1}{p_2} = \frac{U_1(c_1^*, c_2^*)}{U_2(c_1^*, c_2^*)} \simeq \frac{U_1(c^*, c^*)}{U_2(c^*, c^*)} \left[\frac{1 + \varepsilon \frac{U_{11}(c^*,c^*)}{U_1(c^*,c^*)} - \frac{U_{12}(c^*,c^*)}{U_1(c^*,c^*)}}{1 - \varepsilon \frac{U_{22}(c^*,c^*)}{U_2(c^*,c^*)} + \frac{U_{12}(c^*,c^*)}{U_2(c^*,c^*)}} \right]$$

and rearranging this equation we get:

$$\frac{p_1}{p_2} \simeq \frac{U_1(c^*, c^*)}{U_2(c^*, c^*)} \left[1 - \varepsilon \cdot \gamma(c^*, c^*)\right] . \tag{6}$$

Substituting equations (1) and (3), we have

$$MRS(c_1^*, c_2^*) = \frac{p_1}{p_2} \simeq \frac{1}{\delta} - \frac{\varepsilon \cdot \gamma(c^*, c^*)}{\delta} \tag{7}$$

where by definition ε is a function of c_1 and c_2. Therefore, to obtain a point estimates of δ and γ we need observations on consumption choices (c_1, c_2) and the associated relative prices p_1/p_2.

3 Empirical illustration

In this section we illustrate how the theoretical results can be used to estimate time preferences, and apply our approach to the data of Andreoni and Sprenger (2012). We first describe their experimental strategy and the resulting dataset. We then present the maximum likelihood objective function associated with the theoretical results presented above. Finally estimation results are reported.

3.1 Experimental data

In the CTB procedure of Andreoni and Sprenger (2012), subjects have to allocate $m = 100$ tokens across two dates (t and $t + \tau$). Tokens are exchanged for money with exchange rates p_t and $p_{t+\tau}$. Experimental payments are then carried out at the specified dates.[3] Relative prices $p_{t+\tau}/p_{t+\tau}$ implicitly determine the interest rate over the period of τ days. Denoting the allocation of token at each date by x_t and $x_{t+\tau}$, we have $c_t = x_t p_t$ and $c_{t+\tau} = x_{t+\tau} p_{t+\tau}$.

The allocation task was completed by 97 participants. For each participant, both the initial payment date t and the delay before the second payment occurs τ were varied systematically. More specifically, we have $t \in \{0, 7, 35\}$ and delays $\tau \in \{35, 70, 98\}$. For each combination of t and τ, subjects completed the token allocation task for five different relative prices. Thus in total each subject made 45 token allocation decisions, yielding 4365 observations on $(p_t, p_{t+\tau})$ and $(c_t, c_{t+\tau})$.

While the participants can select any allocation of tokens within the interior of their budget, results from the experiment indicate that, for a large number of choice occasions, the selected allocation lie on the boundary. In total around 60 percent of choices are corner solutions, and 36 of the 97 subjects always selected a corner outcome. This suggests high substitutability between payments in the two dates (i.e. allocations display an extreme responsiveness to changes in relative prices), and it also makes the treatment of corner solutions particularly important for

[3] An important feature of the work by Andreoni and Sprenger (2012) is that potential variations in transaction cost associated with different payment dates are minimized through a carefully crafted experimental procedure. For example, thank-you payment was split in two and paid on the dates with planed experimental payments.

the empirical analysis. We come back to this below.

3.2 Empirical estimation strategy

Given $i = 1, \ldots, I$ observations on relative prices (p_1/p_2) and consumption pairs (c_1, c_2), we can write Equation (7) as follows:

$$\left(\frac{p_1}{p_2}\right)_i = \theta_1^{\tau} + \theta_2 \varepsilon_i + u_i \,, \tag{8}$$

where $\varepsilon = (c_t - c_{t+\tau})/2$, t is the number of days before the first payment is carried out, τ is the delay (also in days) between $t = 1$ and $t = 2$, θ_1 and θ_2 are parameters to be estimated from the data, and u_i is an idiosyncratic disturbance term. In principle, this equation can be estimated with OLS to recover an estimate of time preferences:

$$\delta = 1/\theta_1 \,, \quad \text{and} \quad \rho^{annual} = \theta_1^{365} - 1 \,. \tag{9}$$

The CAIS is given by

$$\gamma = -\theta_2/\theta_1 \,. \tag{10}$$

A simple regression analysis thus identifies the MRS on the 45 degree line, together with a measure of the curvature of the intertemporal utility function, without making any assumption about the functional form for $U(c_1, c_2)$.

While this strategy can be straightforwardly applied to interior allocation choices, when a corner outcome is selected the price ratio at the observed consumption choice does not measure the MRS. In particular, when $c_1^* = 0$ then $\mathit{MRS}(0, c_2^*) = p_1/p_2 \geq U_1/U_2$, and when $c_2^* = 0$ is selected $\mathit{MRS}(c_1^*, 0) \leq U_1/U_2$. Thus for corner outcomes, the price ratio represents either a lower bound or an upper bound to the true MRS depending on the specific corner selected.

Importantly, however, corner choices still provide information about bounds of the MRS. To exploit these data points, we treat these as censored observations and add more structure to the error term of the model. Specifically, assuming that u_i is normally distributed with mean zero and variance σ^2, we have that:[4]

$$Prob_i[c_1^* = 0] \;=\; Prob_i \left[MRS_i(c_1^*, c_2^*) < \left(\frac{p_1}{p_2}\right)_i \middle| c_1^*, c_2^* \right] \;=\; \Phi\left[\left\{ \left(\frac{p_1}{p_2}\right)_i - \theta_1^\tau - \theta_2 \varepsilon_i \right\} / \sigma \right] \quad (11)$$

where $\Phi(\cdot)$ is the standard cumulative normal density function. For the second possible corner solution, $c_2^* = 0$, we have:

$$Prob_i[c_2^* = 0] \;=\; \Phi\left[- \left\{ \left(\frac{p_1}{p_2}\right)_i - \theta_1^\tau - \theta_2 \varepsilon_i \right\} / \sigma \right] . \quad (12)$$

Denoting the indicator function by $\mathbb{1}_{c_1^*, c_2^*}$, the log-likelihood function is given by:

$$\begin{aligned} l_i[(p_1/p_2)_i | c_1^*, c_2^*; \theta_1, \theta_2, \sigma] \;=\; & \mathbb{1}_{[c_1^*=0, c_2^*>0]} \log\left\{ \Phi\left[\left\{ \left(\frac{p_1}{p_2}\right)_i - \theta_1^\tau - \theta_2 \varepsilon_i \right\} / \sigma \right] \right\} + \\ & \mathbb{1}_{[c_1^*>0, c_2^*>0]} \log\left\{ \frac{1}{\sigma} \phi \left[\left\{ \left(\frac{p_1}{p_2}\right)_i - \theta_1^\tau - \theta_2 \varepsilon_i \right\} / \sigma \right] \right\} + \\ & \mathbb{1}_{[c_1^*>0, c_2^*=0]} \log\left\{ \Phi\left[- \left\{ \left(\frac{p_1}{p_2}\right)_i - \theta_1^\tau - \theta_2 \varepsilon_i \right\} / \sigma \right] \right\} \end{aligned} \quad (13)$$

where ϕ is the standard normal density function. This expression for the log-likelihood in fact coincides with that of a non-linear two-limit tobit model.[5]

We further consider two extensions to this simple setup. The first quantifies how the delay τ affects the quantities of interest. Intuitively, it may affect the MRS and thus the (annualized) discount rate, as well as the CAIS since increasing the delay may make payments less substitute with each others. Moreover, the analysis of Echenique et al. (2015) suggests that observed choices for different delays are not consistent with exponential discounting. One simple and flexible way to account for the role of varying delays is to estimate equation (7) for different τ separately.

The second extension we consider is the possibility that γ varies with c^*. This involves

[4] Of course, more flexible assumptions can also be considered. For the present analysis, however, we retain the more standard specification, mainly because it is also applied in the original analysis of Andreoni and Sprenger (2012).

[5] We are interested in the (sometimes unobserved) MRS, which is equivalent to the latent variable commonly used to describe tobit-type models. Therefore the estimated parameters readily capture the marginal effects of interest.

specifying γ as a function of c^*, and for simplicity here we assume it is linear: $\gamma(c^*, c^*) = \Gamma_1 + \Gamma_2 c^*$. Hence equation (7) can be rewritten as:

$$\left(\frac{p_1}{p_2}\right)_i = \theta_1^\tau + \theta_2\varepsilon_i + \theta_3\varepsilon_i c^* + u_i\,, \tag{14}$$

where $\Gamma_1 = \theta_2/\theta_1$ and $\Gamma_2 = \theta_3/\theta_1$. This provides additional flexibility in how the curvature of the intertemporal utility function is measured, and therefore potentially a more defendable identification of time preferences (Andersen et al., 2008).

3.3 Results

We now use data by Andreoni and Sprenger (2012) on $p_\tau/p_{\tau+k}$ and consumption choices c_t and $c_{t+\tau}$ in order to compute c^* and ε. We then estimate the parameters in equations (8) to (14) using a maximum likelihood estimator define in Equation (13).

Table 1 reports estimation results for each time delay length and a constant γ (columns 1-3) as well as when γ is a linear function of c^* (columns 4-6). In the table, we first report the parameter estimates (θ's and σ). In the mid-panel, we then provide the implied measures of time preferences (the discount factor δ and the annualized discount rate ρ^{annual}). We also report implied estimates for the parameters of (14) and the implied CAIS, our measure of utility curvature, evaluated at the average constant consumption path. Finally, in the bottom panel of the table, we report the number of observations that represent an interior solution to the utility maximization problem, the number of corner solutions, and AIC/BIC measures of the goodness-of-fit.

Considering first results for (8), reported in columns 1-3, parameter estimates at the top of the table all have the expected sign and are highly statistically significant. This suggests that the experimental procedures generates sufficient variation in allocations to identify key marginal properties of preferences, which can in turn be exploited to estimate time preferences. Ensuing estimates of time preferences presented in the middle panel are precisely estimated, and these suggest an annualized discount rate of around 220% for the shorter payment delay (τ=35) and between 90% and 100% for the longer ones ($\tau \in \{70, 98\}$). Conversely the curvature

Table 1: Results from the empirical illustration

	Constant CAIS (equations 8)			Linear CAIS (equation 14)		
	(1) $\tau = 35$	(2) $\tau = 70$	(3) $\tau = 98$	(4) $\tau = 35$	(5) $\tau = 70$	(6) $\tau = 98$
θ_1	1.003***	1.002***	1.002***	1.003***	1.002***	1.002***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
θ_2	-0.033***	-0.034***	-0.087***	-0.053***	-0.047***	-0.248***
	(0.002)	(0.002)	(0.005)	(0.006)	(0.008)	(0.018)
θ_3	–	–	–	0.002***	0.001*	0.016***
				(0.001)	(0.001)	(0.002)
σ^2	0.149***	0.155***	0.386***	0.147***	0.154***	0.346***
	(0.006)	(0.007)	(0.018)	(0.006)	(0.007)	(0.019)
δ	0.997***	0.998***	0.998***	0.997***	0.998***	0.998***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
ρ^{annual}	2.191***	0.888***	0.894***	2.221***	0.906***	0.974***
	(0.355)	(0.111)	(0.171)	(0.384)	(0.116)	(0.207)
Γ_1	–	–	–	0.053***	0.047***	0.247***
				(0.006)	(0.007)	(0.018)
Γ_2	–	–	–	-0.002***	-0.001*	-0.016***
				(0.001)	(0.001)	(0.002)
$\gamma(\bar{c}^*, \bar{c}^*)$	0.033***	0.033***	0.087***	0.034***	0.033***	0.085***
	(0.002)	(0.002)	(0.005)	(0.002)	(0.002)	(0.005)
N	1455	1455	1455	1455	1455	1455
$N_{c_1>0,c_2>0}$	465	449	415	465	449	415
$N_{c_1=0,c_2>0}$	194	322	231	194	322	231
$N_{c_1>0,c_2=0}$	796	684	809	796	684	809
pseudo-LL	18.55	-12.41	-406.4	22.38	-11.09	-364.1
AIC	-31.09	30.82	818.9	-36.77	30.19	736.1
BIC	-15.24	46.67	834.7	-15.64	51.32	757.2

Notes: Estimation by maximum likelihood with standard errors clustered at the individual-level (reported in parenthesis). ***, **, *: statistically significant at 1, 5 and 10 percent respectively.

parameter γ increases with delay τ, indicating that intertemporal substitutability decreases with the distance between time periods. According to (4) this corresponds to a value for the IES of 12 for $\tau \in \{35, 70\}$ and around 5 for τ=98.

Turning to the case where γ is a function of c^* (equation 14), reported in columns 4-6, we find that both the constant Γ_1 and the slope coefficient Γ_2 are statistically significantly different from zero. The additional flexibility is associated with an improvement of the goodness-of-fit measures (both AIC and BIC decline). Moreover, while the additional flexibility does not impact the estimate of the curvature parameter γ evaluated at the mean of the sample, we find that it implies slightly larger estimates for the discount rates, especially for the longer delays ($\tau \in \{70, 98\}$).

4 Discussion and conclusion

This paper has presented an approach to empirically identify the rate of time preferences based on a simple approximation obtained through the first order conditions of the intertemporal consumption allocation problem. This approach does not require assumptions about the functional form of the utility function, and is therefore non-parametric in the sense that it does not constrain the utility function used to identify time preferences. Another important feature of our approach is that we do not use risk preferences to estimate time preferences.

How do our results compare to those obtained under the assumption of a parametric utility function? In the original analysis of Andreoni and Sprenger (2012), the estimation of time preferences relies on a standard additive intertemporal utility function:

$$U(c_t, c_{t+\tau}) \,=\, u(c_t) + \beta^\tau u(c_{t+\tau})\,,$$

where per-period utility is specified either as as CRRA ($u(c_t) \,=\, \frac{1}{\alpha}(c_t)^\alpha$) or constant absolute risk aversion (CARA, $u(c_t) \,=\, -\exp(-\lambda c_t)$). Aggregate analysis of observed choices using the CRRA utility function suggests an annual discount rate of 0.371 when censoring is not taken into account, 0.324 when censoring is taken into account.[6] Results with a CARA utility function display somewhat lower discount rate, between 0.254 and 0.335 depending on the specification.

Therefore, relative to the results reported in Andreoni and Sprenger (2012), our estimates of annualized discount rates are significantly higher (but not inconsistent with empirical evidence reported elsewhere, see e.g. Frederick et al., 2002, for an early overview). Our results can at

[6] In the Andreoni and Sprenger (2012), this correspond to the baseline case in which background consumption is set to zero. Note that given our definition of the time preferences as a marginal rate of substitution, the choice of non-zero background consumption implies that that β in the additive intertemporal utility function no longer measures time preference (i.e. the expression for the MRS involves more terms). Note also that Andreoni and Sprenger (2012) find no evidence of present bias, and hence we have not considered this possibility here. See also Andersen et al. (2014) for further evidence in this direction.

least partly be explained by the fact that our measure of utility curvature is larger than those reported in Andreoni and Sprenger (2012). This suggest that assuming a constant elasticity property for the per-period utility affects the estimation of the curvature. In turn, as pointed out by Andersen et al. (2008), a lower measure of the curvature increases the implied discount rates identified from observed choices. Another important thing to note is that the functional form assumption used in Andreoni and Sprenger (2012) requires that the elasticity of substitution is independent of the time interval between consumption, so that substitutability of consumption between t and $t + 1$ is the same as between t and $t + \tau$, $\forall \tau$. Our results instead suggest that substitutability varies with both the delay length and the level of consumption.

But beyond specific empirical results, our work suggests a number of avenues for further research on empirical identification of time preferences. In particular, the arguments developed emphasize that experimental procedures designed to measure time preferences should minimize the need for extrapolation. By clarifying dimensions along which this extrapolation has to be undertaken, either using a specific functional form for the utility function or, as we have proposed, with an approximation procedure, our work suggests possibilities to mitigate the associated errors. For example, one possibility is to directly measure relative prices that induces subjects to allocate equal consumption amounts at two different dates. Indeed, based on the the basic definition of the discount factor, this would directly reveal time preferences without the need to extrapolate. Finally, our results show that substitutability of consumption across time is not necessarily constant, and presumably varies both with consumption levels and with delay length. This suggests that further research on the identification of this quantity is a worthwhile endeavor in its own right.

Appendix Coefficient of absolute substitutability

In this Appendix we show that the coefficient of absolute intertemporal substitutability is a property of preferences rather than of a specific utility function.

Consider an increasing function Φ applied to the generic intertemporal utility function $U(c_1, c_2)$. We define

$$V(c_1, c_2) = \Phi(U(c_1, c_2))$$

Denoting the first and second partial derivatives of Φ by Φ' and Φ'' respectively, we have that

$$V_1 = U_1 \Phi' \, , \quad V_2 = U_2 \Phi' \, ,$$

$$V_{11} = U_{11} \Phi' + U_1^2 \Phi'' \, , \quad V_{22} = U_{22} \Phi' + U_2^2 \Phi'' \, ,$$

and

$$V_{12} = U_{12} \Phi' + U_1 U_2 \Phi'' \, .$$

By Definition 3, we have

$$\gamma(c,t)_V \; = \; \left[-\frac{V_{11}}{V_1} + V_{12} \left(V_1^{-1} + V_2^{-1} \right) - \frac{V_{22}}{V_2} \right]^{-1} \, .$$

The different ratios can be computed individually to yield:

$$\frac{V_{11}}{V_1} \; = \; \frac{U_{11}}{U_1} + U_1 \frac{\Phi''}{\Phi'}$$

$$\frac{V_{22}}{V_2} \; = \; \frac{U_{22}}{U_2} + U_2 \frac{\Phi''}{\Phi'}$$

$$\frac{V_{12}}{V_1} \; = \; \frac{U_{12}}{U_1} + U_2 \frac{\Phi''}{\Phi'}$$

$$\frac{V_{12}}{V_2} \; = \; \frac{U_{12}}{U_2} + U_1 \frac{\Phi''}{\Phi'}$$

Collecting the terms we finally get to:

$$1/\gamma(c,t)_V \;=\; \left[-\frac{U_{11}}{U_1} + U_{12}\left(U_1^{-1} + U_2^{-1}\right) - \frac{U_{22}}{U_2}\right] \;=\; 1/\gamma(c,t)$$

and we conclude that γ is a property of the preference relation.

References

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

——— (2011) "Multiattribute utility theory, intertemporal utility and correlation aversion." Working paper 2011-04.

——— (2014) "Discounting behaviour: A reconsideration," *European Economic Review*, 71, pp. 15 – 33.

Andreoni, J. and C. Sprenger (2012) "Estimating time preferences from convex budgets," *American Economic Review*, 102(7), pp. 3357–3376.

Echenique, F., T. Imai, and K. Saito (2015) "Testable implications of models of intertemporal choice: Exponential discounting and its generalizations." Working paper.

Epstein, L. G. (1987) "A simple dynamic general equilibrium model," *Journal of Economic Theory*, 41, pp. 68–95.

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

McFadden, D. (1963) "Constant elasticity of substitution production functions," *The Review of Economic Studies*, 30 (2), pp. 73 – 83.