Time Discounting and Time Consistency

Nicola Dimitri and Jan van Eijck

Abstract Time discounting is the phenomenon that a desired result in the future is perceived as less valuable than the same result now. Economic theories can take this psychological fact into account in several ways. In the economic literature the most widely used type of additive time discounting is *exponential discounting*. In exponential discounting, the fall of valuation depends by a constant factor on the length of the delay period. It is well known, however, that exponential time discounting often does not describe well how people actually behave. Most people are averse to short delays in gratification now, while their future selves may not mind a bit of extra waiting. This behaviour can be described well by non-exponential discounting functions such as *hyperbolic discounting*. In hyperbolic discounting, valuations fall rapidly for small delay periods, but the fall gets slower for longer delay periods. Hyperbolic discounting captures phenomena such as procrastination, addiction and in general inconsistency over time. This chapter investigates whether forms of non-exponential discounting, in particular close to the so called Quasi-Hyperbolic model, could also be characterized in terms of dynamically consistent choices when individuals discount the welfare of future selves as well as their payoffs.

1 Background

When given the choice between receiving one hundred euros today and receiving one hundred euros tomorrow, almost everyone prefers to pocket the cash today. This is an example of a psychological phenomenon that is called time discounting, the phenomenon that a desirable result in the future is perceived as less valuable than the same result now. This all has to do with the "time value of money or goods". A

Nicola Dimitri
Faculty of Economics, Università di Siena, e-mail: `dimitri@unisi.it`
Jan van Eijck
Centrum Wiskunde & Informatica (CWI), Amsterdam, e-mail: `jve@cwi.nl`

sum of money now is worth more than the same sum of money in a year's time. The future value of a payment of $X > 0$ is less than the present value of the payment of X. But *why* is this so?

It is remarkable how little agreement there is on what is at the basis of this phenomenon. In neoclassical economics time discounting is taken as a fact of human nature, and this fact is used to explain the phenomenon of *interest* on capital. The human inclination is to spend now, so the story goes, and the fact that we get a promise of being able to spend *more* later is needed to talk us into foregoing immediate gratification.

Another rather different view, from the Austrian School in economics, is that in a growing economy there can always be a reasonable expectation that supplies will be larger in the future. Thus, consumption of X now is worth more than consumption of X in the future, because one might reasonably expect to be able to consume $X + Y$ in the future. A variation on this would be the expectation of technological progress. Why not buy a smartphone today? Because the same money will buy you a much more advanced model in a year's time.

These two views do not take yet another factor into account: the reasonable expectation that consumption in the future will fail to benefit me in the long run. I might fall ill, or I might even die. What promises of future consumption are worth then is described in this parody of a Salvation Army hymn:

> You will eat, bye and bye,
> In that glorious land above the sky;
> Work and pray, live on hay,
> You'll get pie in the sky when you die.

In any case, accepting a promise of future reward always involves a risk, for promises may not be kept.

As an example of the difficulty of assessing the time value of a good, consider a man of sixty-five, with a well-stocked wine cellar. He would like to be able to drink as long as he enjoys wine, of course. But this does not tell him at which rate to drink. Should he assume he has ten more years to enjoy wine? Twenty more years? Thirty more years? Suppose he decides to be on the safe side, and drink with moderation, so that his stock will last thirty years. Then, after fifteen years, he dies, leaving half a cellar of splendid wines behind to be sold off in auction, with revenues going to the state. What a waste!

Still, one reason to postpone consumption into the future is illustrated by the wine cellar example: diminishing marginal utility. There is a maximum to the number of bottles you can drink in a single day, so even if you do not care about the future you will have to postpone some of your consumption.

Suppose you have a deposit for amount X with your bank. How much should they pay you in n years, just to compensate for the risk that they go bust? Suppose you know what the risk of default per year is. Let's assume that the risk of default on a payment of amount X, say within the next year, is p. Then lending out that amount can be viewed as a bet. You will get X plus interest from the bank with probability $1 - p$. Call the interest I. Since there is a chance of $1 - p$ of receiving $X + I$, a fair bet entails $X = (1 - p)(X + I)$. From this, it follows that $I = \frac{p}{1-p}X$.

To argue why this is reasonable, assume you are taking insurance against the risk of your bank defaulting. What kind of insurance do you need? What should you be willing to pay for the insurance? Since you put X at risk with your bank, your insurance premium I' should yield $X + I'$ in case the bank defaults (the money you lose in the bank crash, plus the premium you paid for insurance), so we get $I' = p(X + I')$. From this we can calculate what the insurance would cost in an arbitrage free market: $I' = \frac{p}{1-p}X$.

Setting $I = I'$, you get a pair of bets without risk, or a so-called *balanced book*: you invest X with your bank and I with your insurance company, and you are sure of getting all your money back: $X + I$ from the bank in case the bank survives, $X + I$ from the insurance company in case the bank defaults. This would be boring, of course, and this is not what people do. What people do is to take risks with money, or arbitrage cheap insurance against risk. But the notion of a reasonable bet still applies, for if the risk is p, for a reasonable bet the total revenue on risking X should be at least $X + I = X + \frac{p}{1-p}X = \frac{1}{1-p}X$. This, of course, does not yet take into account the reward for the temporary transfer of utility of the capital. Conversely, for a chance of $1 - p$ of receiving Y in the future, it is always unreasonable to invest more than $(1 - p)Y$.

For example, if your initial amount is 100 euros, then this is what your bank should pay you, at least, after $0, 1, 2, \ldots$ years, to compensate for a default risk of 10 percent ($\frac{1}{10}$) per year:

$$[100.00, 111.11, 123.46, 137.17, 152.42, 169.35, 188.17, 209.08, 232.31, \ldots]$$

So 100 euros now should be worth 169.35 euros in five years.

Conversely, 100 euros in five years time is worth only 59.05 euros now. The depreciation of an amount of money X in one year, given a risk of default of p is given by $(1 - p)X$. So if the amount is 100 euros, then risk discounting for a stable risk of $\frac{1}{10}$ per year that your bank defaults gives:

$$[100.00, 90.00, 81.00, 72.90, 65.61, 59.05, 53.14, 47.83, 43.05, 38.74 \ldots]$$

To see why risk discounting is an example of what is called exponential discounting, notice that this sequence is calculated by means of:

$$\theta^0 X, \theta^1 X, \theta^2 X, \theta^3 X, \ldots,$$

where $\theta = 1 - p$ is the discount factor.

Now here is a remarkable phenomenon. Given a choice between 100 euros now and 200 euros in a year's time, many respondents prefer 100 euros now. But given a choice between 100 euros in five years and 200 euros in six years, most people report that they prefer the latter. This shows that their discounting factor is not constant. Preference for 100 euros now over 200 euros in a year means that the discount factor *now* is less than $\frac{1}{2}$, for if 100 euros now are preferred to 200 euros in a year it means that $100 > d \times 200$, and therefore $d < \frac{1}{2}$. Preference for 200 euros in six years over 100 euros in five years means that the discount factor *in five years* is more

than $\frac{1}{2}$. Put differently, the assumptions underlying the model of exponential time discounting are questionable.

An excellent introduction to the problems that the exponential time discounting model is facing is given in [6], where the so-called Discounted Utility model proposed by Samuelson [16] is described, and where attention is drawn to a central assumption of this model: all of the disparate motives relevant to intertemporal choice can be condensed into the (constant) discount rate. What constant discounting means is that if a person prefers X at t to Y at $t + d$ for some t, then this same preference must hold for *any t*.

Following the philosopher Derik Parfit, it has become customary to describe the future behaviour of an agent as the behaviour of a 'future self' of the agent [12]. Parfit's philosophy criticises systems of ethics that are based on a notion of personal identity, indeed, he considers the notion of a person at one point in time being identical to the same person thirty years later deeply problematical. Instead, he argues for a view where a person's abuse of tobacco is an action that harms 'another' person, namely that person's future self. Fitting this to the wine cellar example: the 65-year old newly-pensioned man is psychologically only weakly connected to the 95-year old greybeard he may become in thirty years time. This psychological disconnection may be another reason for time discounting, by the way.

The well-known theory of 'rational addiction' of Becker and Murphy [4] assumes that individuals, even when addicted, are time consistent, and that addiction consists in consistently optimizing according to stable preferences. In this model, addicts recognise the full price of their addictive consumption, including the cost of future consumption. From the 'rational addiction' model it follows that all potential harm from addictive behaviour is external. If the decision to indulge in a smoking habit is rational for an individual smoker, a reasonable policy would be to prevent harm to others, and leave it at that.

There is strong empirical evidence, however, that addicted people are unable to carry out stated desires, e.g., the desire to quit smoking, which is at odds with the 'rational addiction' model. When the Becker/Murphy framework is combined with the assumption of hyperbolic discounting, we get a model where addictions which are harmful to the individual are recognized as irrational by the addict. Now smoking has both an external and an internal cost: it causes harm to others but it also harms oneself, and this is recognized by the smoker. So, as Gruber and Köszegi observe in [8], we get radically different recommendations for policy-making.

A converse to addiction is ability to delay gratification, which turns out to be an excellent predictor for success in life, as Walter Mischel at Stanford University found out with his "marshmallow experiment" [11], made popular in Daniel Goleman's bestseller *Emotional Intelligence: Why it Can Matter More than IQ* [7].

2 Recent Developments

The issue of time discounting has recently received increasing attention by a growing body of literature; excellent overviews on the matter are [3], [6], and [10]. A main reason in support of such an effort comes from empirical evidence [1, 2] indicating how intertemporal choices could be better explained by hyperbolic discounting (HD) rather than, the more widely used, exponential discounting (ED). In particular, a version of HD that gained relevance among economists is the Quasi-Hyperbolic Discounting (QHD) proposed in [13, 9]. Indeed, forms of HD seem to provide a better model than ED to explain phenomena exhibiting time inconsistent behaviour like, for example, procrastination.

In fact it is well understood that when an individual's welfare is calculated by an additive functional discounting future (instantaneous utilities) payoffs, time consistent choices, for all payoff streams, can only be the case with ED. Broadly speaking, a preference ordering expressed at time $t = 0, 1, 2 \ldots$ on choices available at time $t^* > t$, will never be changed at time t', with $t \le t' \le t^*$, and $t^* = 1, 2, \ldots$

In this chapter we introduce an alternative, though related, view on intertemporal decision making and show how 'forms' of dynamic consistency are satisfied by generalised versions of ED, that can be close to the most common form of QHD, when an individual is altruistic towards future selves and, as well as her own payoffs, discounts their welfare [15]. Then, a main message of the chapter is that forms of time consistency are not necessarily a prerogative of standard ED. Such a characterisation can also be seen to provide possible foundations for more general models than ED. The result is elicited from a benchmark environment given by an infinite stream of unitary payoffs, available to an individual over time. The rationale behind it is that a stationary framework should make the essential properties of intertemporal preferences emerge in a more natural way.

The main finding of this work, which confines itself to the analysis of a specific type of altruism towards future selves, establishes that when time consistency holds, purely egoistic behaviour and the mixed egoistic-altruistic type of behaviour that we investigate imply the same individual welfare, for a given degree of altruism.

3 The Model and Main Results

Consider an individual evaluating an infinite sequence x of constant numbers $\{x(t)\}$, with $x(t) = c$ and $t = 0, 1, 2, \ldots$, where the index t is interpreted as time. To keep the exposition simple, and with no loss of generality, let $c = 1$. As it is standard, we conceive such an individual as a collection of selves, one for each t.

In this section of the chapter we shall be studying the following additive functional form:

$$W(s; T, K) = \sum_{t=s}^{s+T-1} d(t-s) + \sum_{t=s+T}^{s+T+K} W(t; T, K)d(t-s), \tag{1}$$

where $W(s; T, K)$ is a finite number representing the welfare of a self s, calculated by discounting payoffs up to time $s + T - 1$ and then the welfare of future selves from time $s + T$ to $s + T + K$, with $s, K = 0, 1, 2 \ldots$ and $T = 1, 2, \ldots$ We interpret K as the degree (extent) of altruism towards future selves, while T determines the first future self to enjoy altruistic concern on the part of self s.

Notice that T should have been written as $T(s)$, since we want to contemplate each self to possibly have a different T; this is why our notion of time consistency will be asked to hold for each T. Finally, $d(0) = 1$ and $0 \leq d(t) \leq 1$, with $t = 1, 2, \ldots$, is a converging sequence of numbers. We refer to the sequence $\{d(t)\}$ as the individual's discounting function. Expression (1) could be given at least two interpretations. The first would simply say that proper welfare calculation is done through recursive relations of some type, rather than by mere payoff discounting, since a rational agent should take decisions considering her welfare at future times. The second, as we mentioned, simply views consideration about one's future welfare as a form of altruism towards future selves [15]. In what follows we shall refer to both readings.

It is worth anticipating that in this chapter the following specific forms of (1) will have a special part.

$$i) \quad W(s; \infty, K) = \lim_{T \to \infty} W(s; T, K) = \lim_{T \to \infty} \sum_{t=s}^{T} d(t - s). \tag{2}$$

This is the most common (standard) criterion of payoff discounting used to calculate an individual's welfare through additive functional forms. It could be thought of as either a completely egoistic (with respect to future selves) approach or else fully myopic.

$$ii) \quad W(s; T, 0) = \sum_{t=s}^{s+T-1} d(t - s) + W(s + T; T, 0)d(T) \text{ with } T < \infty. \tag{3}$$

In this simplest mixed type of discounting, the summation considers payoffs up to some time $s + T - 1$ and then only the welfare of self $s + T$, where the expression is truncated.

$$iii) \quad W(s; T, \infty) = \lim_{K \to \infty} W(s; T, K) = \tag{4}$$

$$\lim_{K \to \infty} \left[\sum_{t=s}^{s+T-1} d(t - s) + \sum_{n=s+T}^{s+T+K} W(t; T, K)d(n - s) \right].$$

In (4) altruism takes its fully extended form in discounting the welfare of all future selves from $s + T$ onward.

We now come to the notion of time consistency. To motivate it we observe how in such a stationary framework, represented by the infinite sequence of unitary payoffs (and for given K), there should be no a priori reason for the selves to perceive different welfare, independently of the value of T. Namely, the notion that we have

in mind is robust with respect to the point in time when self s starts being altruistic, and so sensitive to K only. A full discussion on the relation between our notion and the standard definition of time consistency is future work (but see Section 4 for some connections to the literature).

The idea is captured by the following definition:

Definition 1 (Time Consistency). We say that an individual's preferences are "Time Consistent" (TC) if, given K, the discounting function $d(t)$ is such that

$$W(s; T, K) = W(s'; T, K),$$

for all $s \neq s' = 0, 1, 2, \ldots$ and all $T = 1, 2, \ldots$.

We are now ready to formulate the main result of the chapter.

Theorem 1. *(a)* $W(s; T, K)$*, with $T < \infty$, $K < \infty$, and $W(s; \infty, K)$, with $K < \infty$, satisfy TC only if*

$$d[t(K + 1) + i] = d(i)D(K)^t \tag{5}$$

for $i = 1, 2, \ldots, K + 1$; $t = 0, 1, 2, \ldots$ and $D(K) < 1$, where $D(K) = \sum_{i=1}^{K+1} d(i)$.
(b) $W(s; T, \infty) = 1$ *for all $T < \infty$.*

Proof. (a) Let $T, K < \infty$. If $W(s; T, K)$ satisfies TC then

$$W(s; T, K) = W = W(s'; T, K)$$

for all $s \neq s' = 0, 1, 2, \ldots, T = 1, 2, \ldots$ and given K.
Let $s = 0$ so that

$$W(0; T, K) = W = d(0) + [d(1) + \cdots + d(K + 1)]W = 1 + D(K)W,$$

from which

$$W = \frac{1}{1 - D(K)} = W(s; T, K).$$

Notice now that

$$W(s; T + 1, K) - W(s; T, K) = W - W = Wd(T + K + 1) + d(T) - Wd(T) = 0,$$

from which

$$d(T + K + 1) = (1 - 1/W)d(T) = D(K)d(T),$$

and it is immediate to verify that the conclusion to the first part follows.
Consider now $T = \infty$. In this case,

$$W(s; \infty, K) = 1 + \sum_{t=0}^{\infty} \left[\sum_{i=1}^{K+1} d(i)D(K)^t \right] = 1 + \sum_{t=1}^{\infty} D(K)^t = \frac{1}{1 - D(K)} \text{ for all } s$$

only if $D(K) < 1$.

(b) Take now $W(s; T-1, \infty) = W(s; T, \infty)$. Hence, $W(s; T, \infty) - W(s; T-1, \infty) = 0$, which leads to $d(T)(1 - W) = 0$, satisfied when $W = 1$ or, equivalently, $d(T) = 0$ for all $T = 1, 2, \ldots$ □

In words, the above theorem says that if the stationarity of the environment is captured, and so TC holds, the discounting function must have a form akin to the standard ED. In particular, this obtains when $K = 0$; indeed, in this case $D(0) = d(1)$ and $d[t(K + 1) + i] = d(t + 1) = d(1)[d(1)^t] = d(1)^{t+1}$ for $t = 0, 1, 2, \ldots$.

This suggests how our notion of time consistency would compare with the standard one. With $K = 0$ we have ED, namely the ratio between two consecutive terms of the discounting function $\frac{d(t+1)}{d(t)} = D(1)$ is constant. If, in the standard notion, this guarantees consistency at each point in time, and for each payoff profile, which individual would $K > 0$ model? It is easy to check that now it is the ratio $\frac{d(t+K)}{d(t)} = D(K)$ of K terms apart of the discounting function that remains constant. This would guarantee time consistency in the standard sense, but only for intervals that are multiples of K. What this means is that, if at date t an individual prefers option A to option B, then at times $t + nK$, with $n = 1, 2, \ldots$ the individual would still prefer A to B.

Consider, for example, a job contract between an employer and a salesman. The contract specifies that if at the end of each year the revenues accruing to the employer from the sales of the employee (say $Y \in$), are above some specific target (say $\overline{Y} \in$), then the salesman would get a bonus of some fixed percentage X of the total sales. The contract also specifies that the employer may reconsider his initial decision every 12 months. If his time preferences are given by the discounting function $d(i)D(K = 12)^t$, with $i = 1, 2, \ldots$, then this would mean the employer never wants to change his initial decision. Therefore, if the employer only reconsiders his decisions at specific times which are 12 months apart from each other, he would be time consistent in the standard sense.

When $K \geq 1$, expression (5) can specify forms close to QHD. In particular, consider the following example.

Example 1. Example Let $d(i) = [d/(K + 1)]$, for $i = 1, \ldots, K + 1$ and $0 < d < 1$. Then $D(K) = d$ and $d[t(K + 1) + i] = [d/(K + 1)](d^t) = [d^{t+1}/(K + 1)]$. Putting $[1/(K + 1)] = \beta < 1$, the discounting function then becomes the sequence

$$1, \beta d, \ldots, \beta d, \beta d^2, \ldots, \beta d^2, \ldots, \beta d^t, \ldots, \beta d^t, \ldots$$

which bears similarities to the standard QHD $1, \beta d, \beta d^2, \ldots, \beta d^t$ function, except that now the generic term βd^t characterises more than one term of the function. The fact that the discounting function remains constant between t and $t + K$, with $t = 1, 2, \ldots$, means that the individual assigns the same importance to his future selves in that cohort ("generation"). So dates $t = nK$, with $n = 1, 2, \ldots$ could be interpreted as time instants where the individual perspective, of the future, changes, after having considered on an equal footing the previous cohort of K selves.

The coefficient β would have a simple explanation, being inversely related to the degree of altruism (number of discounted future selves' utilities). The larger

K, the lower the weight associated to each discounting coefficient; hence, the more concerned about the welfare of future selves an individual is, the lower her own welfare. At the limit, when $K \to \infty$ this would tend to 1.

In words, the above result provides the following, specific, connotation of the generalised forms of QHD identified by (5). In the stationary environment that we investigate, they represent the only type of intertemporal preferences that would cause completely "egoistic" individuals (as described by 2), and "altruistic" individuals (for a finite degree K of altruism) to perceive the same welfare within (namely independently of T, for T finite) and between the two criteria.

4 Discussion

In this chapter we characterized generalized forms of Quasi-Hyperbolic Discounting as the time consistent preferences of an individual evaluating a stationary environment, represented by an infinite stream of unitary payoffs available over time. In particular, the result obtains when an individual discounts the welfare of her future selves as well as her future payoffs; namely, when she has a certain degree of altruism. Such connotation can also be seen as suggesting a way to find a foundation for additive time preferences different from Exponential Discounting.

Rohde [14] proposes a very simple way to calculate the measure of time inconsistency from two data: indifference between reward X at t and Y at $t + s$, and indifference between X at $t + u$ and Y at time $t + u + v$. In future work, we hope to connect this work to the findings above.

A rather different explanation of hyperbolic discounting is proposed by Peter Sozou in [17]. This gives a mathematical formalization of why hyperbolic discounting is rational for an individual that considers his own hazard of survival as uncertain. This applies to the case of the sixty-five year old wine cellar owner. Using a mortality table he can fix the 'correct' discount rate using $1 - q_t$, for the average chance of surviving his next birthday at age t. But the problem is that this is an average, and there is uncertainty as to how well this average matches the survival chance of the actual individual. Here is a relevant quote from the Sozou paper:

> Some authors have interpreted this time-preference reversal effect as indicating non-rational time preferences [...]. Thus I may appear to be temporally inconsistent if, for example, I prefer the promise of a bottle of wine in three months over the promise of a cake in two months, but I prefer a cake immediately over a promise of a bottle of wine in one month.
>
> There is, however, no inconsistency if I perceive a promised future reward not as a sure thing, but instead as having a probability attached to it. This can be illustrated numerically. [...]

This is on a very different tack from the approach taken in Section 3: instead of 'explaining away' time inconsistency, as we have arguably done, Sozou explains its rationale. Clearly, if Sozou is right, then it seems there is no need for trying to invoke notions of dynamic consistency. It would be very interesting to do some further investigations to compare the two views.

When is it rational to prefer a cake now over a promise of a bottle of wine in a month, and also to prefer a bottle of wine in three months over a cake in two months? This is rational in case the promise of the wine in a month looks dubious now, but less dubious in two months from now. The decision maker takes the growth in confidence about the probability of the rewards into account in her decision making.

On the other hand, an account of decision making in terms of uncertainties about risk would call for prudence in proportion to our uncertainties about the risks we are running. This is not what we actually observe. A report that there is (a bit) *less certainty* about the causes of climate change than the warnings of the Intergovernmental Panel on Climate Change (IPCC) have made us believe until now, is taken by many people as a reason to worry *less* about the prospect of climate change. If Sozou is right, people should worry *more*. But if people sometimes discount less in contexts with more uncertainty, one would like a model that explains *that* as well (see also [5]).

References

1. AINSLIE, G. *Picoeconomics*. Cambridge University Press, Cambridge, 1992.
2. AINSLIE, G. *Breakdown of Will*. Cambridge University Press, Cambridge, 2002.
3. ANGELETOS, G., LAIBSON, D., REPETTO, A., TOBACMAN, J., AND WEINBERG, S. The hyperbolic consumption model: Calibration, simulation, and empirical evaluation. *Journal of Economic Perspectives 15* (2001), 47–68.
4. BECKER, G. S., AND MURPHY, K. M. A theory of rational addiction. *Journal of Political Economy XCVI* (1988), 675–700.
5. DASGUPTA, P., AND MASKIN, E. Uncertainty and hyperbolic discounting. *American Economic Review 95*, 4 (September 2005), 1290–1299.
6. FREDERICK, S.AND LOEWENSTEIN, G., AND O'DONOGHUE, T. Time discounting and time preferences: A critical review. *Journal of Economic Literature 40* (2002), 351–401.
7. GOLEMAN, D. *Emotional Intelligence: Why it Can Matter More than IQ*. Bantam, 1995.
8. GRUBER, J., AND KÖSZEGI, B. Is addiction "rational"? theory and evidence. *The Quarterly Journal of Economics* (2001), 1261–1303.
9. LAIBSON, D. Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics 112* (1997), 443–477.
10. LAIBSON, D. *Encyclopedia of Cognitive Sciences*. Nature Publishing Group, London, 2003, ch. Intertemporal decision making.
11. MISCHEL, W., SHODA, Y., AND RODRIGUEZ, M. L. Delay of gratification in children. *Science 244* (1989), 933–938.
12. PARFIT, D. *Reasons and Persons*. Oxford University Press, 1986.
13. PHELPS, E. S., AND POLLACK, R. On second-best national national saving and game-equilibrium growth. *Review of Economic Studies 35* (1968), 185–199.
14. ROHDE, K. I. M. The hyperbolic factor: A measure of time inconsistency. *Journal of Risk and Uncertainty 41* (2010), 125–140.

15. SAEZ MARTI, M., AND WEIBULL, J. Discounting and altruism to future decision-makers. *Journal of Economic Theory 122* (2005), 254–266.

16. SAMUELSON, P. A. A note on measurement of utility. *The Review of Economic Studies 4*, 2 (1937), 155–161.

17. SOZOU, P. D. Hyperbolic discounting and uncertain hazard rates. *Proceedings of the Royal Society, London 265* (1998), 2015–2020.