

Choosing Monetary Sequences: Theory and Experimental Evidence

Paola Manzini

Marco Mariotti

Luigi Mittone

University of St. Andrews and IZA

University of St. Andrews

University of Trento

This version: November 2009

Abstract

We formulate and investigate experimentally a model of how individuals choose between time sequences of monetary outcomes. The model assumes that a decision-maker uses, sequentially, two criteria to screen options. Each criterion only permits a decision between some pairs of options, while the other options are incomparable according to that criterion. When the first criterion is not decisive, the decision maker resorts to the second criterion to select an alternative. We find that: 1) traditional economic models based on discounting alone cannot explain a significant (almost 30%) proportion of the data no matter how much variability in the discount functions is allowed; 2) our model, despite considering only a specific (exponential) form of discounting, can explain the data much better solely thanks to the use of the secondary criterion; 3) our model explains certain specific patterns in the choices of the ‘irrational’ people. We reject the hypothesis that anomalous behaviour is due simply to random ‘mistakes’ around the basic predictions of discounting theories: deviations are not random and there are clear systematic patterns of association between ‘irrational’ choices.

J.E.L.codes: C91, D9

Keywords: time preference, time sequences, negative discounting.

Corresponding address: School of Economics and Finance, University of St. Andrews, Castlecliffe, The Scores, St. Andrews KY16 9AL, Scotland, UK. Part of this work was carried out while Manzini and Mariotti were visiting the University of Trento, and funding for the experiments was provided by the ESRC under grant RES-000-22-0866. We wish to thank both institutions for their support. We are also grateful to the thoughtful comments of two referees and to Alan Agresti, Glenn Harrison, John Hey, Stepana Lazarova, Daniel Read, Bob Sugden and seminar audiences at the LSE Choice Group, the 2006 FUR Conference and the University of Birmingham for insightful discussions and comments, as well as to the tireless staff of the CEEL lab in Trento, in particular to Marco Tecilla for superb programming support. All errors are our own.

1 Introduction

Most economic decisions involve a time dimension, hence the need for a reliable model of time preferences. The standard exponential discounting model for time preferences has been the object of strong, evidence-led criticisms in the last few years. Various discounting ‘anomalies’ have been identified. Some of these anomalies do not contradict the basic maximisation hypothesis in economic theory, and can be addressed simply by changes in the functional form of the objective function which agents are supposed to maximize.¹

But other observed violations of the standard model are more fundamental, because they seem to contradict the basic assumption of maximization of *any* economically reasonable objective function. One of these hard anomalies is the striking phenomenon of *negative time preferences*. Notably Loewenstein and Prelec [10] (also Loewenstein and Prelec [12]) have argued that there is evidence of negative time preferences when individuals choose between *sequences* of outcomes (e.g. wage profiles in a survey by Loewenstein and Sicherman [13] and discomfort sequences in Varey and Kahneman [25]). Their empirical findings lead them to conclude that ‘To most persons, a deteriorating series of utility levels is a rather close approximation to the *least* attractive of all possible patterns’ (p. 347). In a more recent survey, Frederick, Loewenstein, O’Donoghue [7] emphasise once again that ‘In studies of discounting that involve choices between two outcomes... positive discounting is the norm. Research examining preferences over *sequences* of outcomes, however, has generally found that people prefer improving sequences to declining sequences’ (p. 363). These findings obviously cannot be explained by hyperbolic discounting, or indeed by any other form of positive discounting. Therefore they pose a more formidable challenge for the economic modeler of decision-making over time sequences.

This paper has three main aims:

1. We propose a theory of preferences over monetary time sequences that provides a possible explanation for the observed anomalies, while at the same time keeping simple

¹The best-known example of a soft anomaly in choice over time is preference reversal between date-outcome pairs, which can be explained by the now popular model of hyperbolic discounting, as well as by other models (see e.g. Manzini and Mariotti [17], Noor [18], Read [19], Rubinstein [21]). Recall that preference reversal denotes the situation whereby the preference for a smaller reward obtained sooner over a larger later reward is reversed when the obtainment of both rewards is pushed forward by the same amount of time. We survey the topic of anomalies in Manzini Mariotti [16].

exponential (positive) discounting as one of its core elements.

2. We investigate to what extent conclusions on the preference for increasingness based on survey findings (e.g. Loewenstein and Sicherman [13]) are supported in a laboratory experiment in which subjects received real money payments – we find mixed evidence on this.
3. We uncover some new clear patterns in choice (beside preference for increasing sequences) which are consistent with our theory but inconsistent with pure discounting models.

While experimental investigations of choices over date-outcome pairs form a small but non-negligible literature, experimental investigations of choices over reward sequences are extremely thin on the ground in the economics literature, especially with financially motivated subjects.² In the experimental part of the paper, we ask subjects to make binary choices among all possible pairs of monetary sequences, out of a set of an increasing, constant, decreasing and ‘jump’ (i.e. end effect) pattern in a paid condition (where subjects do indeed receive the sums corresponding to the sequence chosen).³

The theoretical model we propose is a ‘hybrid model’ which combines the traditional consistent discounting theory with a heuristics component. More specifically, in order to rank monetary reward sequences, the decision maker looks first at the standard exponential discounting criterion; however, preferences are *incomplete*, so that sequences are only *partially* ordered by the criterion. In other words, sometimes the decision maker is able to make a trade-off between the time and the outcome dimension, and sometimes he is not. When he is, he does so in a time-consistent way.⁴

We postulate a very simple (two-parameter, in the specification used for the experiment) *interval order* structure to formalize preference incompleteness. In this structure, preferences

²The literature includes Chapman [4], Gigliotti and Sopher [8], Guyse, Keller and Epple [9]. The last two papers in particular show that the domain of choice is rather important, in the sense that there are differences in observed choices depending on whether or not the sequences are of money, or health or environmental outcomes. On this point see also Read and Powell [20].

³As a robustness check, we also run additional treatments varying the amount of the payments.

⁴This idea originated in Manzini and Mariotti [17]’s model of choice between date/outcome pairs. We showed that in that context a simple model can account for major anomalies both soft and hard, such as preference reversal and cyclical choice patterns.

are described by two functions, a *utility function* u and a *vagueness function* σ , which combine additively. A sequence a is definitely chosen over another sequence b if the discounted utility of a exceeds the discounted utility of b by at least $\sigma(b)$. When sequences cannot be compared by means of discounted utilities, the decision-maker is assumed to focus on one prominent attribute of the sequences. This prominent attribute ranks (maybe partially) the sequences and allows a specific choice to be made. This latter aspect of the model is in the spirit of Tversky, Sattath and Slovic [24]’s *prominence hypothesis*. The attribute may be context dependent. In the date-outcome pairs case, for example, objects have two obvious attributes that may become prominent, the date and the outcome: in that case two natural models emerge according to whether date or outcome is looked at first. We stress that, at the abstract level, the only departure from the standard choice theoretic approach is that our decision maker’s behavior is described by combining sequentially two possibly incomplete preference orderings, as in Manzini and Mariotti [14], instead of using directly a complete preference ordering.

In the case of reward *sequences*, previous experimental evidence suggests that the general trend of the sequence (increasing or decreasing) is relevant to make decisions. However, in our case the data provide much weaker evidence than Loewenstein and Prelec [10]’s in support of their view that ‘sequences of outcomes that decline in value are greatly disliked’ (p. 351). We find that, even in the simple decision problems we study, where monetary sequences can be clearly ordered according to their trends, simply choosing according to the heuristics that favors an increasing trend, though compatible with a non negligible proportion of choices, does a rather poor job at explaining the data. The modal subject and choice is ‘rational’, in the sense of being compatible with positive time preference combined with preference for income smoothing (concave utility function).

So although there *is* a problem for pure discounting standard theory, its magnitude is not of the scale the existing literature suggests. When there are no affective factors involved (such as, for example, the sense of dread for choices relating to health, or the sense of failure involved in a decreasing wage profile), *some* theory of positive discounting can provide a rough approximation of the choice patterns.

Nonetheless, it is still true that a disturbingly high number of people (around 30%) choose

in ways that are incompatible with *any form* of positive discounting (exponential, hyperbolic or otherwise). This proportion of people violating such a basic economic assumption (that good things should come early and bad things should come late) is unsatisfactory from the point of view of the descriptive adequacy of standard theory. It suggests that other mechanisms beyond discounting are at work. So we believe that Loewenstein and Prelec’s pioneering findings do capture, beside affective factors, some of the heuristic considerations that people use when evaluating ‘neutrally’ (without affects) money sequences. However, those considerations become really effective in explaining the deviant choices only when used as a ‘secondary criterion’, rather than directly. The very basic two-parameter version of our model is far superior, in order to explain observed choices, both to any pure discounting model and to a direct heuristics-based model. In addition, when specialized to date-outcome pairs comparisons, it can also explain other observed anomalies both soft and hard.

Most importantly, while our model nests the standard exponential discounting model as a special case (when vagueness is sufficiently small), it does *not* include the hyperbolic discounting model as a special case. Nevertheless, our theory is in principle able to accommodate more patterns of choice: so of course one would expect a more general theory to be able to explain more data. In order to address this issue, we use Selten’s [23] *measure of predictive success*, and show that our model performs also *proportionally* better than any generic discounting model (including hyperbolic discounting), no matter the degree of concavity of the utility function. We obtain these results by comparing the explanatory power of two alternative *classes* of theoretical models: that of vague time preferences and that of standard discounting models. It is this comparison, rather than the estimation of discount factors,⁵ that we seek to address in this paper.

One important feature of our analysis is that we delve quite deeply into the analysis of ‘irrational’ choices. First of all, a caveat. We use the terms ‘rational’ and ‘irrational’ by implicitly identifying a monetary sequence with the *consumption* sequence associated with it by immediately spending the money when it becomes available. Otherwise, one could not

⁵As explained very well in Andersen, Harrison, Lau and Rutstrom [3], the careful estimation of discount factors would require the simultaneous estimation of risk aversion, in turn requiring to make specific assumptions on the functional forms for decision makers’ utility functions. As we seek to compare large classes of competing models, this goes counter to narrowing down to a specific functional form for the utility function, thus making it impossible to estimate discount factors.

justify as ‘rational’ even the choice of a constant sequence over a decreasing sequence (any consumption smoothing afforded by a constant sequence is also afforded by a decreasing sequence with the same total value). So this terminology, which should be viewed as merely conventional rather than substantial, errs on the side of caution in identifying ‘irrational’ sequences.

We find that the observed pattern of irrationality is systematic. In general, our data reveal some interesting and non-obvious patterns of association in choice, on which standard theory (and simple increasingness heuristics) are completely silent. Among our findings are the following two: (i) *there is association between certain types of rational choices and irrational choices* (those who prefer a decreasing to a constant sequence are disproportionately concentrated among those who also prefer a constant to an increasing sequence); (ii) *there is association between irrational choices of a different type* (choosing an increasing over a decreasing sequence is very strongly associated with choosing an increasing over a constant sequence). Such patterns are what one would expect if our model were true. They cannot be generated by any discounting model.

2 A Model of Intertemporal Choice

Let X indicate a set of money amounts and $u : X \rightarrow \mathcal{R}$ be an instantaneous monotonic increasing utility function. Let $T = \{0, 1, 2, \dots, T^*\}$ be a finite set of dates. The set of alternatives A is a subset of the set of finite sequences of outcomes, i.e. $A \subseteq X^T$. A typical alternative is denoted $a = ((a_1, t_1); \dots; (a_k, t_k))$, with $a_1, \dots, a_k \in X$ and $t_1, \dots, t_k \in T$, and $t_i > t_{i'}$ for $i > i'$.

Recall that in discounting models sequences are evaluated by means of a discounting function $\delta : T \rightarrow (0, 1)$. The discounted utility at time 0 of sequence $a = ((a_1, t_1); \dots; (a_k, t_k))$ is $\sum_{i=1}^k \delta(t_i) u(a_i)$; with exponential discounting we have that $\delta(t) = \delta^t$ for some $\delta \in (0, 1)$. With hyperbolic discounting instead $\delta(t)$ is a hyperbolic function of the type $(1 + at)^{(-g/a)}$, with g and a being two preference parameters, and in the very popular $(\beta - \delta)$ version of hyperbolic discounting the following specification is used: $\delta(0) = 1$ and $\delta(t) = \beta\delta^t$ for $t > 0$, with $\beta, \delta \in (0, 1)$ (see Loewenstein and Prelec [11]).

The model we propose uses as a primitive a binary preference relation \succsim^* of the individual on A constructed as follows:

1. *Primary criterion.* There exists a primary criterion P_1 , which is a possibly incomplete strict ordering. On the basis of the primary criterion the individual makes (possibly partial) comparisons between sequences.
2. *Vagueness function.* There may be pairs of alternatives for which the primary criterion alone is not discriminating enough. This lack of discrimination is captured by a ‘vagueness function’ $\sigma : A \rightarrow \mathcal{R}_+$.
3. *Secondary criterion.* In the case where the primary criterion does not rank alternatives (i.e. the decision maker is *vague*), and only in this case, a secondary criterion is used. The secondary criterion P_2 is just a (possibly partial) strict ordering on A .

We interpret P_1 as resolving comparisons for which the trade-off between outcomes and time yields, in the perception of the individual, a decisive advantage to one of the alternatives. The trade-offs involved are assumed to be resolved by P_1 in a ‘time-consistent’ way: they coincide with the standard ones based on exponentially discounted utility, with discount factor $\delta \in (0, 1)$.⁶ If the present, exponentially discounted utility of the higher value alternative a does not exceed the utility of the lower value alternative b by at least $\sigma(b)$, we say that the decision maker is *vague*. In this case a and b are not related by P_1 . Then the decision between a and b is resolved based on the secondary criterion P_2 . We interpret it as being based on one prominent attribute of the elements of A . In some contexts, as in the case of date-outcome pairs or in the experiment presented below, the relevant attributes are obvious; in other cases less so and the issue of what is an appropriate secondary criterion is essentially empirical. We view the secondary criterion P_2 as a primitive of the model just as u , β or δ are primitives in the simple version of the hyperbolic discounting model. We build on the empirical evidence discussed previously and posit that the preference for increasing sequences kicks only when resolving the vagueness of the decision maker; so we adopt a preference for increasingness as our secondary criterion.

⁶As noted in the conclusions, other assumptions on P_1 are also compatible with the data.

The (complete) relation \succsim^* is derived by the combination of the (possibly incomplete) primary and secondary criterion, as follows. Let \succ^* denote the strict binary preference relation on A . We propose the following general model, for given u, δ, σ and P_2 :

For all $a, b \in A$, we have $a \succ^* b \Leftrightarrow$

1. $\sum_t \delta^t u(a_t) > \sum_t \delta^t u(b_t) + \sigma(b_t)$ (Primary Criterion P_1), or
2. $\sum_t \delta^t u(a_t) \leq \sum_t \delta^t u(b_t) + \sigma(b_t)$ and $a P_2 b$ (Secondary Criterion P_2)
 $\sum_t \delta^t u(b_t) \leq \sum_t \delta^t u(a_t) + \sigma(a_t)$

(where the summations are taken over the appropriate range).

The model we propose here is grounded in previous work: the existence of two criteria P_1 and P_2 , applied sequentially to arrive at a choice can be justified at a more abstract level - Manzini and Mariotti [14] provide an axiomatic foundation for an abstract choice function (taken as a primitive) to be ‘rationalisable’ by a two-stage procedure of the type specified in this paper.⁷ The specialisation of P_1 that we propose in this paper has quite conveniently an additive form in which the vagueness term enters the formula. A few manipulations show that this specialization can account for both cyclical behaviour in choice as well as other patterns that cannot be accommodated in any discounting model.⁸ We pursue these points further in section 4.

3 Experimental Design

Our objective in this experiment is to compare the explanatory power of alternative *classes* of theoretical models: we do not wish to pin down any particular functional form, as this would open up issues as to the suitability of the chosen specification for the problem at hand. For example, if one wanted to consider a concave discounted utility function, why might a Cobb-Douglas be more appropriate than a quasi-linear specification? Our approach frees us

⁷In Manzini and Mariotti [17] we provide a characterisation for the $(\sigma - \delta)$ model used in this paper.

⁸Rubinstein [22] reports that subjects exhibited the following type of behavior: they chose \$997 to be received on November 1st over \$1000 to be received on December 1st (they were impatient and preferred smaller reward earlier rather than larger reward later) but chose the sequence of four payments of \$1000 each to be paid on the first day of April, June, October and December over the sequence of four payments of \$997 each available one month earlier. This choice pattern, incompatible with discounting models, can be accommodated within the $(\sigma - \delta)$ model (see http://www.st-andrews.ac.uk/~pm210/expsequencestechnical_and_data_appendix.pdf).

from the need to estimate any functional form parameters, and allows us to proceed along a different path.⁹ We will consider the large family of standard pure discounting theories where the utility over date-outcome pairs is separable in the time and outcome components (i.e. for outcome a_i available at time t_i , discounted utility is $\delta(t_i) u(a_i)$), and the utility for a stream of date-outcome pairs is additively separable across periods (i.e. given by $\sum_{t=t_1}^{t_k} \delta(t) u(a_t)$), with the usual monotonicity and concavity properties. Similarly for the class of vagueness models. In the experiment, we elicit choices between pairs of alternative remuneration plans; given all possible combinations, and regardless of specific functional forms, as we detail below, there will be only a subset of all possible choice patterns compatible with the family of pure discounting theories. Similarly, only some of all possible patterns are compatible with our vagueness model. Thus if the standard model is to be successful in explaining the data, there must be only a handful of observed choice behaviours incompatible with such theoretical framework. Similarly, a measure of success of our proposed alternative model would have to show that a consistent proportion of the data falls in categories that are not compatible with standard pure discounting theories, but allowable in the new model. In addition, any successful model would have to be falsifiable, i.e. at the theoretical level there would have to be choice patterns incompatible with it, giving evidence a chance to contradict the theory.

One potential difficulty is that the framework of vague time preferences does include standard discounting as a particular case. As a consequence, it would not be surprising if the vagueness family were to perform better than pure discounting theories at explaining the data. To control for this fact, in comparing theories we will rely on Selten's index of explanatory power, devised precisely to deal with such instances. As we will see, based on this index one can proclaim a theory as more successful than another only if it accommodates *proportionally* (as opposed as raw data) more data than the competitor.

In our approach, then, rather than pursuing the (point) estimation of specific functional form parameters, we investigate the compatibility of our two competing 'area theories', or rather of the subsets of the universe of possible choice patterns that they can accommodate, with observed data. In this sense, we proceed in the spirit of traditional revealed preference

⁹See Abdellaoui, Attema and Bleichrodt [1] (and the literature therein) for a recent experimental estimation of discount rates, which finds little evidence for hyperbolic discounting.

analysis. An additional advantage of this line of attack is also that we do not have to impose on our experimental subjects a long list of questions (which would be needed if we wanted to estimate, say, individual discount rates), as in effect only a limited number will be enough for us to be able to make meaningful comparisons between competing families of theories. These questions are detailed in the next sections.

Experiment

The experiments were carried out using the facilities of the Computable and Experimental Economics Laboratory at the University of Trento, in Italy. In all, we ran 16 sessions. Experimental subjects were recruited through bulletin board advertising from the students of the University of Trento. Each sessions consisted of both male and female participants in roughly equal proportions. The experiment was computerised, and each participant was seated at an individual computer station, using separators so that subjects could not see the choices made by other participants. Experimental sessions lasted an average of around 26 minutes, of which an average of 18 minutes of effective play, with the shortest one lasting approximately 16 minutes and the longest around 37 minutes. In our main treatment, the **PAY** (for Paid) treatment, subjects were paid €5 showup fee and an additional €48 paid in instalments depending on their choice (a total of 102 subjects in 9 sessions). In addition, we ran two additional treatment as a robustness check. In the **HYP** (for hypothetical) treatment subjects were only paid the €5 show up fee, whereas in the **PAY_L** (for Paid, low stakes) the total additional amount on top of the show up fee was €24.¹⁰ These were carried out to check whether the amounts paid would make a difference to observed choice behaviour: a comparison between **PAY** and **PAY_L** would allow us to check whether the size of the real monetary payments produces any effects; while a comparison between **PAY** and **HYP** would allow us to check whether there is any difference between real and hypothetical payments. Since however the results we obtain in these additional treatment are qualitatively similar to those for the main treatment, they are not reported here.¹¹ At the beginning of the

¹⁰In all treatments the show up fee alone, for an average of less than thirty minutes long experimental session, was higher than the hourly pay on campus, which is €8. At the time of the experiments the exchange rate of the Euro was approximately €1=\$1.2=£0.7. In terms of purchasing power €1 was approximately equivalent to £1.

¹¹Available online at http://www.st-andrews.ac.uk/~pm210/expsequencestechnical_and_data_appendix.pdf.

experiment a screenshot with instructions appeared on each monitor, and at the same time an experimenter read the instructions aloud to the participants.¹² In each treatment, each experimental subject was presented with 23 different screens. Each screen asked the subject to choose the preferred one among a set of alternative remuneration plans in installments to be received staggered over a time horizon of nine months, each consisting of €48 overall (€24 in the **PAY_L** treatment). It was explained that at the end of the experiment one screen would be selected at random, and the preferred plan for that screen would be delivered to the subjects.¹³ As usual, subjects were free to abandon the experiment at any point in time (although no one did). The payment was not directly managed by the experimenters, but through the accounts office of the University of Trento, where subjects would pick up their payment at the established times (and they would receive several reminders from the lab before each date. All subjects have been paid). This procedure, which was explained to the participants before the beginning of the experiment, was followed in order not only to preserve anonymity with respect to the experimenters, but to provide a further guarantee that payments will actually be forthcoming.¹⁴

Choices were based on two sets - depending on the number of installments - of four plans each, namely an increasing (I), a decreasing (D), a constant (K) and a jump (J) series of payments, over either two or three installments, as shown in Table 1.¹⁵ Though in both cases payments extended over nine months, because of the different number of installments we abuse terminology and refer to ‘two-*period*’ (or also ‘short’) sequences and ‘three-*period*’

¹²The translation of the original instructions (in Italian) can be found in the appendix.

¹³Instructions were the same in all treatments, bar for one sentence, which in the HYP treatment clarified that choices were purely hypothetical, so that the only payment to be received would be the show up fee.

¹⁴This payment procedure is customarily used in the CEEL lab to pay participant. It is very difficult to rely on alternative payment methods, given that many students do not have current accounts and at any rate the banking system in Italy is still highly bureaucratic. Using other instruments such as Paypal is more costly (in Italy there is a charge to receive money, at least there was at the time of the experiment) and anyhow electronic money transfer are highly disliked, as internet transactions are still not trusted in general, and even more so at the time of the experiment. On the other hand, as the experimental lab has a long tradition, we do not believe that issues of trust in receiving delayed payments were relevant. Observe further that even if one were to assume that subjects did not trust us to pay them, all sequences have a front end delay, so that there is no reasons to expect any one sequence to be preferred to any other on the basis of subjects mistrust. Furthermore, if mistrust influenced choice, there should be no differences between the PAY and the HYP treatment, whereas we find to the contrary.

¹⁵The corresponding reward schedules for the **PAY_L** treatment are obtained by dividing each amount by two (e.g. the K2 sequence consisted of two equal payments of €12 each, the K3 sequence of three equal payments of €8 each, and so on).

(or also ‘long’) sequences rather than two/three-*installment* sequences.

Pairwise choices, which are the interest of this paper, were interspersed with choices between larger sets of sequences (analysed elsewhere¹⁶). In fact each subject had to make a selection from each possible subset of plans within each group (making up 11 choices per group).

Two period sequences					Three period sequences				
	I2	D2	K2	J2	I3	D3	K3	J3	
in three months	16	32	24	8	8	24	16	8	in three months
					16	16	16	8	in six months
in nine months	32	16	24	40	24	8	16	32	in nine months

Table 1: the base remuneration plans

Participants made their choice by clicking with their mouse on the button corresponding to the preferred remuneration plan. Once made, each choice had to be confirmed explicitly, so as to minimize the possibility of errors. Both the order in which the questions appeared on screen and the position of each plan on the screen was randomised. Sample screenshots are in figure 1.




Figure 1: Sample screenshots

¹⁶See Manzini and Mariotti [15].

4 The theoretical models and the experimental framework

To make life harder for our proposed model, in this section we consider a simplified version of the ‘vagueness’ model where we constrain the σ function to be just a constant. We refer to this two-parameter version as the ‘ $(\sigma - \delta)$ model’. Note that this restriction if anything limits the ability of our model to fit the data. Before deriving its predictions for choice in our experimental setup, though, we review briefly what the predictions of standard discounted utility theories are in this context, since we will be comparing precisely these two classes of theories. To distinguish them from the $(\sigma - \delta)$ model, in which the primary criterion is also based on discounting, we will refer to the large family of standard discounted utility theories (which includes e.g. hyperbolic discounting) as *pure discounting theories*. As above, we assume the utility over date-outcome pairs to be separable in the time and outcome components, that is $U(a, t) = \delta(t) u(a)$, and let the utility for a stream of date-outcome pairs be given by $U((a_1, t_1), (a_2, t_2), \dots, (a_k, t_k)) = \sum_{i=1}^k \delta(t_i) u(a_i)$, i.e. additively separable. Let u be monotonic increasing in outcome, concave and with positive third derivative.¹⁷ Finally, let the discounting function be monotonically non increasing, i.e. $t > t' \Rightarrow \delta(t) \leq \delta(t')$.

Now fix the times at which outcomes are received as 0, 1, 2, so that sequences can be defined in terms of the ordered outcomes with no mention of the time at which each reward is available. To reduce notation let $u_1 = u(8)$, $u_2 = u(16)$, $u_3 = u(24)$ and $u_4 = u(32)$; and for the discounting function let $\delta(0) = \delta_0$, $\delta(1) = \delta_1$, and $\delta(2) = \delta_2$. Also, normalize the utility function u (by dividing it by δ_0) so that we can set $\delta_0 = 1$. One important caveat: the above notation is assuming implicitly that the monetary amounts accrued to the experimental subjects as immediately consumed, and hence consider utility for money directly (as opposed to mediated by consumption). A full discussion of this assumption is deferred to section 6.

Let \succ_d denote the preference relation of a decision maker who discounts utility available at time t by some discount function (i.e. a ‘pure discounter’). With the four three period

¹⁷Note that these are mild assumptions, satisfied by the most common functional forms used in economics to describe an agent’s utility function, such as for instance the constant risk aversion class of utility functions.

sequences described in the right hand panel of Table 1, any discounting criterion for choice \succ_d should order them as either $D \succ_d K \succ_d I \succ_d J$ or $K \succ_d D \succ_d I \succ_d J$. The choice $I \succ_d J$ holds since J shifts some outcome from the second to the last period while increasing the sequence dispersion, so that any discounting criterion, paired with the concavity of the utility function, is going to favor I over J . The choice between D and K also depends on the shape of the utility function. We indicate each sequence by the letter and the number of installments in which it was paid; for example $J3$ refers to the three period jump sequence (when we do not want to emphasize the length, we just use the letters). Then, straightforward manipulations yield that D is chosen over K whenever the discount factor is small enough.¹⁸

Secondly, regardless of sequence length, both D and K must be preferred to I , and the latter must be preferred to J . Consequently the patterns of choice which can be observed when a decision maker has utility for monetary streams which are additively separable and who has a non-decreasing discount function are either of the following two:

- $D \succ_d K \succ_d I \succ_d J$, or
- $K \succ_d D \succ_d I \succ_d J$.

This is true even if the preferences of the decision maker conform to hyperbolic discounting, since what matters is only the assumption that the discount function is monotonically nondecreasing. For instance, in the case of the $(\beta - \delta)$ model we would have $\delta_1 = \beta\delta$ and $\delta_2 = \beta\delta^2$, which would not affect the analysis above.

Consider now the $(\sigma - \delta)$ model. The primary criterion compares the present discounted utility of monetary streams. As explained before, we impose strong conditions on the discount function by letting $\delta_t = \delta^t$. We assume that the secondary criterion orders by *increasingness*. For the simple sequences of payments K , I and D listed above, if a decision maker is vague between any two sequences, by the secondary criterion it must be that

$$I \succ_2 K \succ_2 D \tag{1}$$

¹⁸Formally, $D3 \succ_d K3 \Leftrightarrow \delta_2 < \frac{u_3 - u_2}{u_2 - u_1} \equiv \delta_L$ and $D2 \succ_d K2 \Leftrightarrow \delta_2 < \frac{u_4 - u_3}{u_3 - u_2} \equiv \delta_S$. All the conditions that follow derive from straightforward manipulations comparing the utilities from each reward sequence. The tedious calculations can be consulted online at the following address: http://www.st-andrews.ac.uk/~pm210/expsequencestechnical_and_data_appendix.pdf

with \succ_2 transitive.

In general, putting both primary and secondary criterion together, easy manipulations show¹⁹ that for the $(\sigma - \delta)$ model, I can be preferred to either K or D provided that the vagueness σ is sufficiently high for the secondary criterion to kick in; similarly for a preference of K over D . Moreover, the parameter values are such that they imply that whenever a decision maker chooses D over K , it **must** be the case that he chooses K over I and D over I . On the other hand, a choice of K over D imposes no such restrictions on the choice between either I and K or I and D . In short, depending on parameter values, the $(\sigma - \delta)$ model can accommodate the following choices:

$$\begin{aligned} I2 \succ^* K2 &\Leftrightarrow \sigma \geq a & I3 \succ^* K3 &\Leftrightarrow \sigma \geq A \\ I2 \succ^* D2 &\Leftrightarrow \sigma \geq b & I3 \succ^* D3 &\Leftrightarrow \sigma \geq B \\ D2 \succ^* K2 &\Leftrightarrow \sigma < c & D3 \succ^* K3 &\Leftrightarrow \sigma < C \end{aligned}$$

where a, A, b, B, c and C are real numbers whose precise values depend on the utility function and on the discount factor.²⁰

As a final point, note that the above rests on the assumption that subjects are risk averse or neutral. This mild assumption is in accordance with previous experimental evidence (see e.g. Andersen, Harrison, Mortensen and Rutström, [3]). Assuming risk loving behaviour would only make the case against standard discounting models stronger, as e.g. no preference for the constant sequence over the decreasing sequence could be explained.

5 Experimental Results

The sample consisted of 102 experimental subjects, roughly in equal proportions across sexes.²¹ To indicate the choice of one plan over another we use the ' \succ ' symbol, e.g. $K2 \succ D2$ indicates that in two periods sequences, the constant one was chosen over the decreasing one. We will use the ' \succ^* ' and ' \succ_d ' notation when discussing the predictions for choice behavior

¹⁹See http://www.st-andrews.ac.uk/~pm210/expsequencetechnical_and_data_appendix.pdf for the full derivation.

²⁰Full details are available at http://www.st-andrews.ac.uk/~pm210/expsequencetechnical_and_data_appendix.pdf

²¹All the exact statistical analysis has been carried out using *StatXact*, v.7. For a thorough reference on exact methods in categorical data analysis see Agresti [2].

according to the $(\sigma - \delta)$ model and all pure discounting theories, respectively.

5.1 Aggregate data

Frequency distributions for binary choices involving the base sequences of payments I , K , J and D are reported in Table 2. They show the following:

	2 periods (%)	3 periods (%)
D chosen over K	66.7	64.7
D chosen over I	79.4	81.4
D chosen over J	90.2	84.3
K chosen over I	92.2	93.1
K chosen over J	89.2	91.2
I chosen over J	92.2	91.2

Table 2: Frequency distribution of binary choice, aggregate data (102 subjects)

1. **Sequence length does not matter:** the only difference in binary choice behavior when moving from two to three period sequences which is statistically significant is for the choice between J and D, where the proportion of subjects preferring J over D increases from 9.8% to 15.7% (the p-value for the corresponding Mc Nemar test is 0.035). This seems to suggest that when the ‘jump’ aspect of the J sequences kicks in (i.e. for the two period sequences J is simply steeper than I), it does affect choice behavior;
2. **A majority of subjects prefers decreasing to increasing sequences:** this is in sharp contrast with the suggested preference for increasing sequences discussed in the introduction;
3. **A majority of subjects prefers rational to irrational sequences:** the constant sequence is preferred to both the increasing and the jump sequence more than 90% of the times, and the decreasing sequence is preferred to both the increasing and the jump ones, though somewhat less decisively (more than 80% of the times). Indeed, regardless of length, the subjects who chose I over D are almost thrice as many as

those choosing I over K²² (the corresponding proportions are 20.6% against 7.8% for the short sequences and 18.6% against 6.9% for the long sequences);

4. **Endpoint effect?** For the long sequences, subjects choosing the jump series over the decreasing one are almost twice as many than those choosing the jump sequence over the constant one²³ (15.7% against 8.8%), whereas for the short sequences the frequency is approximately the same (recall that for short sequences, J and I are in fact both increasing, with the J sequence steeper than the I one. For long sequences, though, the end effect in the J sequence comes to the fore).

These data already suggest that decision making is unlikely to be guided by a clear-cut discounted utility rule: choice of either the increasing or jump sequence over either the constant or the decreasing one is a sign of ‘irrationality’, so that any individual choosing the increasing sequence over either the constant or the decreasing sequence displays choice behavior which is incompatible with pure discounting theories.

5.2 Checking theories

There are eight possible profiles of choice generated by the three binary comparisons involving the I, K and D sequences:

Profile 1: $D \succ K \succ I$;

Profile 2: $K \succ D \succ I$;

Profile 3: $D \succ I \succ K \succ D$ (cycle);²⁴

Profile 4: $I \succ K \succ D$;

Profile 5: $K \succ I \succ D$;

Profile 6: $D \succ K \succ I \succ D$ (cycle);

Profile 7: $I \succ D \succ K$;

²²This difference is statistically significant: a McNemar test of the difference between the proportion of subjects choosing I over K and those choosing I over D returns a p-value of 0.001 for the short and 0.002 for the long sequences.

²³This difference is statistically significant: a McNemar test of the difference between the proportion of subjects choosing K over J and those choosing D over J returns a p value of 0.046.

²⁴Our observation that the proportion of subjects choosing I over D is much higher than the proportion of subjects choosing I over K seems to indicate that $A > B$, so that the bottom portion of Figure 2 should be the one that applies, i.e. we should expect to find no subjects whose choices conform to profile 3. One can proceed similarly for the two period profiles, exchanging a for A , b for B and c for C .

Profile 8: $D \succ I \succ K$.

The first two profiles are the only admissible ones in any model of pure discounting. As for the $(\sigma - \delta)$ model, recall from section 4 that:

$$\begin{aligned} I2 \succ^* K2 &\Leftrightarrow \sigma \geq a & I3 \succ^* K3 &\Leftrightarrow \sigma \geq A \\ I2 \succ^* D2 &\Leftrightarrow \sigma \geq b & I3 \succ^* D3 &\Leftrightarrow \sigma \geq B \\ D2 \succ^* K2 &\Leftrightarrow \sigma < c & D3 \succ^* K3 &\Leftrightarrow \sigma < C \end{aligned}$$

Since it is always the case that $c < a, b$ and $C < A, B$, so that, as we mentioned already, whenever a decision maker chooses D over K, he must also choose K over I and D over I. On the other hand, a choice of K over D imposes no such restrictions on the choice between either I and K or I and D.

In short, there are only five (out of the eight possible) preference profiles which are compatible with the $(\sigma - \delta)$ model, as in Figure 2.




Figure 2: Admissible choice profiles in the $(\sigma - \delta)$ model

The observed frequency distribution of the choice profiles is reported in table 3. It shows that there is a substantial proportion of subjects (around 18%) whose choice cannot be accounted for by standard discounting theories, no matter how flexible the functional form, but that can be accommodated within the $(\sigma - \delta)$ model. That is, looking at the profiles which are not common to the two approaches, there are substantially more observations in

profiles 3-5 than in profiles 6-8 (18 cases in profiles 3-5 as opposed to 5 or 3 cases in profiles 6-8, depending on sequence length), suggesting at first glance that standard discounting fails proportionally more often than the $(\sigma - \delta)$ model in explaining the data. As we will see, this informal observation is confirmed by more careful statistical analysis below.

Choice profiles	Two period sequences (%)	Three period sequences (%)
1: $D \succ K \succ I$	61.77	61.77
2: $K \succ D \succ I$	15.69	17.65
3: $(D \succ I \succ K)$	0.98	0.98
4: $I \succ K \succ D$	5.88	4.9
5: $K \succ I \succ D$	10.78	11.76
6: $(D \succ K \succ I)$	3.92	1.96
7: $I \succ D \succ K$	0	0
8: $D \succ I \succ K$	0.98	0.98

Table 3: frequency distribution of choice profiles for two and three period sequences, PAY treatment (102 subjects)

Since for each subjects we observed choices in *two* alternative settings (i.e. long and short sequences), it makes sense to check how many *subjects* made choices that conform to the two broad families of theoretical explanations. to do so we must first examine the relationship that the two families of models postulate between choice profiles in the two and three period cases. In fact, the choice profile for sequences of a given length may determine the choice profile for the sequences of other length.²⁵ In particular, for the $(\sigma - \delta)$ model the juxtaposition of the two graphs for the choice profiles in the case of two and three period sequences reveals that a switch either from $D \succ^* K \succ^* I$ to $K \succ^* D \succ^* I$ or the opposite switch from $K \succ^* D \succ^* I$ to $D \succ^* K \succ^* I$ with sequence length is possible. This point is visualized in Figure 3.

In figure 3 we denote by δ_i , with $i = L, S$, the critical value of the discount factor such that $c, C > 0$. Then D is chosen over K in a i sequence (i.e. short or long) of payment provided that σ is sufficiently small (i.e. $\sigma < c$ and $\sigma < C$, respectively). Since $\delta_L \equiv \frac{u_3 - u_2}{u_2 - u_1}$ and $\delta_S \equiv \frac{u_4 - u_3}{u_3 - u_2}$, our assumptions on the concavity of the utility function imply that $\delta_L < \delta_S$ always. In the top panel of Figure 3, which applies whenever $\delta^2 \in (0, \delta_L)$, a $(\sigma - \delta)$ decision maker with $\sigma > c$ will exhibit the choice profile $K \succ^* D \succ^* I$ when choosing among

²⁵This follows from straightforward though lengthy and tedious calculations, available at http://www.st-andrews.ac.uk/~pm210/expsequencestechnical_and_data_appendix.pdf

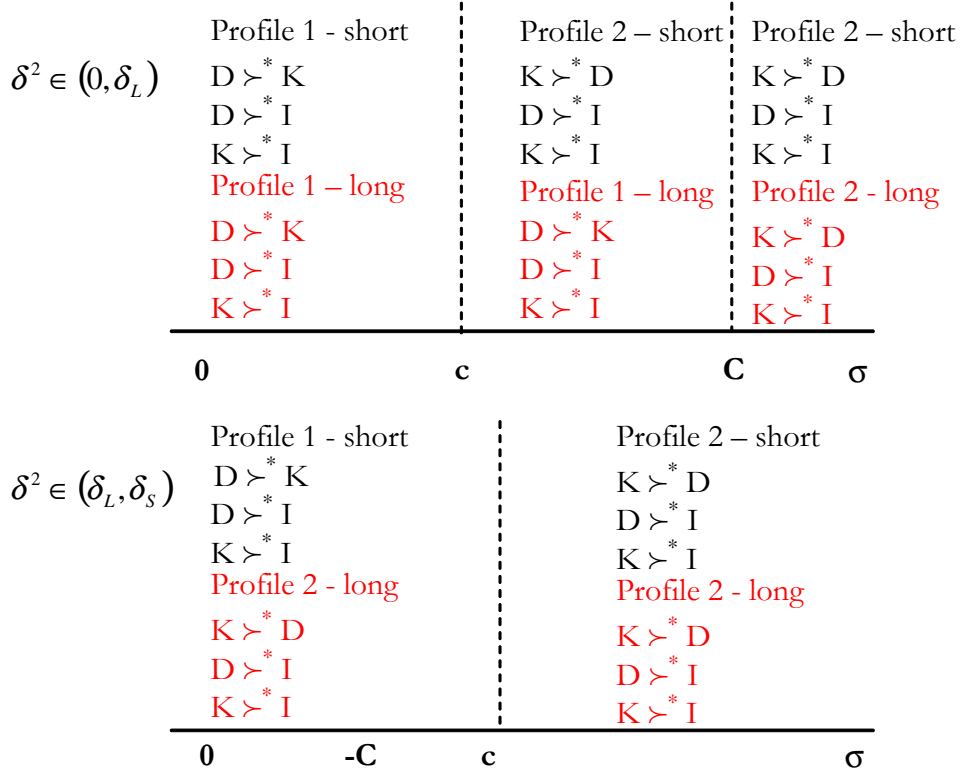


Figure 3: Choice profiles over two and three period sequences compatible with the $(\sigma - \delta)$ model

two period sequences, and the profile $D \succ^* K \succ^* I$ when choosing among three period sequences. An opposite switch will be displayed by an individual whose discount factor is such that $\delta^2 \in (\delta_L, \delta_S)$, and for whom the value of $\sigma < c$.

In contrast, only one of these switches is admissible according to pure discounting theories. Recall that D is chosen over K if and only if the two period discount function is sufficiently small, with the smaller threshold δ_L applying to the case of three period sequences, and the larger threshold δ_S applying to two period sequences. We show this in Figure 4 below, only the switch from $D \succ_d K \succ_d I$ to $K \succ_d D \succ_d I$ is possible when increasing sequence length.

Remarkably, then, despite the fact that the $(\sigma - \delta)$ model is restricted to exponential discounting, in this experimental setup it necessarily explains more choice profiles than *any* pure discounting theory allowed to use any form of discounting, including hyperbolic discounting.²⁶

²⁶Obviously this is not the case in general.

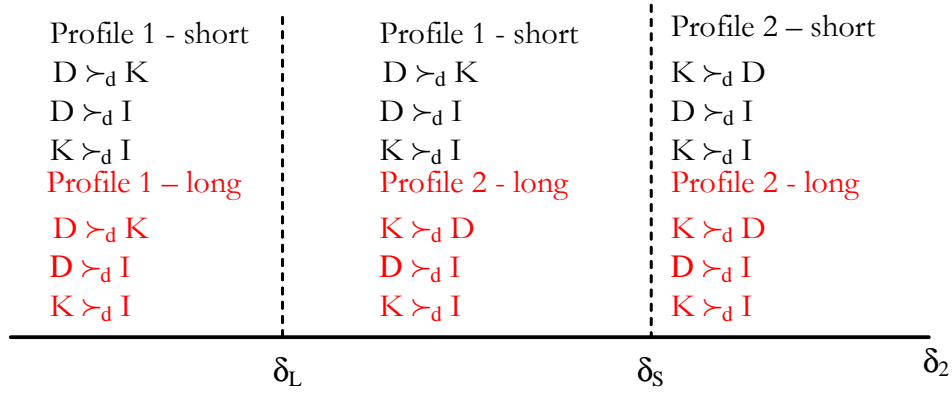


Figure 4: Choice profiles admissible in discounting models

Turning now to the data, table 4 displays the cross-tabulation of the choice profiles observed in two period (columns) and three period (rows) sequences. Number in parentheses refer to the overall percentage of cases. For legibility, diagonal observations (where no change in choice profiles is observed with sequence length) are in bold; groups of subjects whose preference profile amount to at least 5% of the total are highlighted in italics. Combinations of choice profiles compatible with the $(\sigma - \delta)$ model are underlined.

3\2	D>K>I	K>D>I	(D>I>K)	I>K>D	K>I>D	(D>K>I)	I>D>K	D>I>K	Total
D>K>I	<i>55 (53.9)</i>	<i>6 (5.9)</i>	0	0	0	1	0	1	63 (61.8)
K>D>I	<i>8 (7.8)</i>	<i>9 (8.8)</i>	<u>1</u>	0	0	0	0	0	18 (17.6)
(D>I>K)	0	0	0	<u>1</u>	0	0	0	0	1 (1)
I>K>D	0	0	0	2	<u>2</u>	1	0	0	5 (4.9)
K>I>D	0	1	0	<u>3</u>	<i>7 (6.9)</i>	1	0	0	12 (11.8)
(D>K>I)	0	0	0	0	1	1	0	0	2 (2)
I>D>K	0	0	0	0	0	0	0	0	0
D>I>K	0	0	0	0	1	0	0	0	1 (1)
Total	63 (61.8)	16 (15.7)	1 (1)	6 (5.9)	11 (10.8)	4 (3.9)	0	1 (1)	102 (100)

Table 4: choice profiles for two and three period sequences, PAY treatment

The 6 subjects whose preferences fall in the (1, 2) cell cannot be accommodated within any pure discounting model, nor can the other 24 subjects whose choice profiles fall anywhere in the table apart from the first 2 by 2 submatrix.

All in all, then, about 30% of subjects (i.e. 30 out of 102) display a pattern of choice

incompatible with *any* discounting model. To the contrary, only 8 subjects' choice behavior is incompatible with the $(\sigma - \delta)$ model.²⁷

We summarize the ability of pure discounting theories and of the $(\sigma - \delta)$ model to explain data in Table 5, where the proportion of subjects whose choices cannot be accounted in standard and non-standard (e.g. hyperbolic) discounted utility frameworks is more than three and a half times than in the $(\sigma - \delta)$ model.

	explained	Unexplained	Total
Any discounting	72 (70.6%)	30 (29.4%)	102 (100%)
$(\sigma - \delta)$ model	94 (92.2%)	8 (7.8%)	102 (100%)

Table 5: explanatory power of competing theories, PAY treatment

Since the two models are nested, in order to compare differences in explanatory power we compute *Selten's index of predictive success* (see Selten [23]) for both of them. This measure allows the evaluation of 'area theories', namely theories that exclude deterministically a subset of the possible outcomes. Selten's index considers both the 'descriptive power' of the model (measured by the proportion of observations consistent with the model being studied) and its 'parsimony' (i.e. the proportion of cases theoretically compatible with the model under consideration): the lower the proportion of theoretically possible outcomes consistent with the model, the more parsimonious the model.

In the case of the $(\sigma - \delta)$ model one possible criticism of might be that the experiment does not have enough power to reject its predictions even if it happened to be the wrong maintained hypothesis. By introducing the 'parsimony' element, Selten's measure would pick up this problem. More precisely, the measure, denoted s , is expressed as

$$s = r - a$$

where r is the descriptive power (number of actually observed outcomes compatible with the model divided by the number of possible outcomes) and a is the 'relative area' of the

²⁷Arguments similar to those used in the main text to explain preference shift with sequence length can be used to show that for the 3 subjects exhibiting the shift from $I2 \succ K2 \succ D2$ to $K3 \succ I3 \succ D3$, the change is compatible with the $(\sigma - \delta)$ model if $\delta^2 \in (\delta_L, \bar{\delta})$ and $\sigma \in (a, A)$. For the 2 subjects with the opposite shift one needs $\delta^2 \in (\bar{\delta}, \bar{\gamma})$ and $\sigma \in (b, A)$. For the subject with profiles $(D2 \succ I2 \succ K2 \succ D2)$ and $K3 \succ D3 \succ I3$ we need $\delta^2 < \delta_L$ and $\sigma \in (a, A)$, and for the subject with choice profiles $I2 \succ K2 \succ D2$ and $(D3 \succ I3 \succ K3 \succ D3)$ we need $\delta^2 < \delta_L$ and $\sigma \in (b, B)$.

model, namely the number of outcomes in principle compatible with the model divided by the number of all possible outcomes. Selten argues that a theory is better than a competing one if it has a higher value of the index s which combines hit rate and parsimony in a linear way. The reason is pretty straightforward: one would expect a theory with lower parsimony to be able to account for more observations, but one should weigh the gain in hit rate against the cost in parsimony. At one extreme, a vacuous theory (i.e. one that cannot be falsified) would have an index of zero, since it has a zero parsimony (i.e. $a = 1$) but also perfect hit rate.

Selten's index of predictive success is particularly useful in the present context, where the $(\sigma - \delta)$ model can in principle explain more patterns of behaviour than the class of pure discounting theories, and so is less parsimonious.

As we saw above, the class of pure discounting theories is compatible with three possible configurations of choice over two and three period sequences out of the possible 84, while there are eleven patterns of choice compatible with the $(\sigma - \delta)$ model. That is, using subscripts d and $(\sigma - \delta)$ for the two theories we have $a_d = \frac{3}{84} = 0.047$ and $a_{(\sigma - \delta)} = \frac{11}{84} = 0.172$. The data in Table 5 provide us with $r_d = 0.706$ and $r_{(\sigma - \delta)} = 0.922$, so that Selten's indices are $s_d = 0.659$ and $s_{(\sigma - \delta)} = 0.750$.

This confirms the superior performance of the simple $(\sigma - \delta)$ model with respect to the whole class of standard discounting theories, even taking into account the parsimony of the model rather than just its hit rate.

The distinctly larger predictive power of the $(\sigma - \delta)$ model, even with restrictions on the form of discounting used as primary criterion, as compared to conventional theories seem to stem from the fact that, as shown in table 3, not all failures of the standard theories (profiles 3-8) are observed in the same measure, with some of them being observed more often than others. To pursue this point further, it is instructive to rearrange Table 3 as Table 6.

Table 6 displays the cross-tabulation of choices involving the two rational sequences on the one hand, and one irrational sequence on the other. In this way we can address the question of whether or not departures from rationality are generated by independent mistakes, e.g. 'trembles' at the moment of making a decision. Inspection of Table 6 makes it clear that the answer is negative, as independence is strongly rejected; Fisher's exact p-value is 0.001

	D2>K2	K2>D2		D2>K2	K2>D2
I2>K2	1	7	I2>D2	4	17
K2>I2	67	27	D2>I2	64	17

	D3>K3	K3>D3		D3>K3	K3>D3
I3>K3	1	6	I3>D3	2	17
K3>I3	65	30	D3>I3	64	19

Table 6: choices with and without ‘irrational’ sequences

(resp. 0.007). Indeed, in the top (respectively, bottom) panel, for the table on the left, the sample odds ratio²⁸ is 0.057 (resp. 0.076). In other words, the odds of being rational by preferring the constant over the increasing sequence of payments when the decreasing sequence is preferred to the constant sequence are over 17 times (resp. 13 times) the odds of being rational when the constant sequence is preferred to the decreasing one).

Observe the particularly counterintuitive nature of this association: the fact that K is chosen against D makes it *less* likely that it will be chosen against I! This would be very hard to explain in any ‘preference ordering plus error’ model even with a special, non-independent error structure. Similar patterns are found for the right-hand tables. In the right table the odds ratio are 0.062 for the two-period sequences and 0.035 for the three period sequences. Again independence is clearly rejected in both cases (Fisher’s exact p-values are 0.005 and 0.003).

In summary, *subjects who make an irrational choice (either I over K or I over D) are disproportionately concentrated among those who prefer the constant to the decreasing sequence.*

According to pure discounting theories, any choice between K and D should give no information about the distribution of the other binary choices. Thus if one were to cross-tabulate the choice between K and D against the other choices, there should be no association. Yet, consider table 7.

The percentage of subjects choosing D over I and I over K, corresponding to the last column in each of the tables above, is tiny (just 2 subjects for both sequence lengths, that is

²⁸Recall that for the table

	B_1	B_2	
A_1	w	x	
A_2	y	z	

the odds ratio is obtained as $\frac{(w/x)}{(y/z)}$.

Two period sequences				
	K2 \succ I2, D2 \succ I2	K2 \succ I2, I2 \succ D2	I2 \succ K2, I2 \succ D2	I2 \succ K2, D2 \succ I2
D2 \succ K2	63	4	0	1
K2 \succ D2	16	11	6	1

Three period sequences				
	K3 \succ I3, D3 \succ I3	K3 \succ I3, I3 \succ D3	I3 \succ K3, I3 \succ D3	I3 \succ K3, D3 \succ I3
D3 \succ K3	63	2	0	1
K3 \succ D3	18	12	5	1

Table 7: rational versus irrational sequences

less than 2%). On the other hand, the percentage of subjects choosing I over D and K over I, corresponding to the second column in each of the tables above, is around 14% (15 and 14 subjects in the two and three period sequences, respectively). Of these, the overwhelming majority (11/15 and 12/14) lies in the second row of each table, i.e. subjects who also prefer the constant to the decreasing sequence. The effect is stronger for the longer sequences (where arguably the ‘sequence’ feature is more apparent).

Each of the cells in the two tables corresponds to one of the eight possible profiles of choice generated by the three binary comparisons involving the I, K and D sequences. Only profiles 1 and 2 are admissible in any model of pure discounting, and correspond to the first column in Table 7 above. Observations in any of the other cells of the columns could only be due to mistakes. Yet, the association between rows and columns in Table 7 is immediately apparent. In fact the Fisher-Freeman-Halton exact test strongly rejects the null hypothesis of independence (the exact p-value is less than 0.001 for both tables). We conclude that no model of positive discounting is compatible with our data.

But as we saw, on the contrary the association in Table 7 is to be expected in the $(\sigma - \delta)$ model, as the choice between D and K *is* informative on the distribution of the other choices.

Conclusions in a similar vein can be drawn by analysing the cross-tabulation of the choices involving one rational and one irrational sequence, as in by considering a Table 8 and 9 below. Again they reject the hypothesis that failures of rationality are due to random mistakes.

Both Tables 8 and 9 display a strong association between choices between I and D and between I and K. The sample odds ratios are 15.8 and 14.46 in Table 8 and 14.3 and 16.6,

Two period sequences			Three period sequences		
	I2 \succ K2	K2 \succ I2		I3 \succ K3	K3 \succ I3
I2 \succ D2	6	15	I3 \succ D3	5	14
D2 \succ I2	2	79	D3 \succ I3	2	81

Table 8: rational versus increasing sequences

Two period sequences			Three period sequences		
	J2 \succ K2	K2 \succ J2		J3 \succ K3	K3 \succ J3
J2 \succ D2	5	5	J3 \succ D3	6	10
D2 \succ J2	6	86	D3 \succ J3	3	83

Table 9: rational versus jump sequences

respectively, in Table 9, with 95% exact confidence intervals whose lower bounds are all above 2: in other words, the odds of choosing an irrational sequence over a rational one are at least twice as large when an irrational choice has been made between a different rational/irrational pair.²⁹ Thus the hypothesis of independence in choices in the two situations can be rejected (Fisher’s exact test yields p-values of 0.001, 0.002, 0.001 and 0.001). This suggests that there is some systematic mechanism generating the irrational choices that makes irrational choices in one context (e.g. I versus D) strongly associated with irrational choices in another context (I versus K). Whatever this mechanism is, as we saw above it makes the proportion of ‘mistakes’ in the I versus D context significantly higher than the proportion of mistakes in the I versus K context (see footnotes 22 and 23).

6 Money, consumption, budget sets and choices

As already mentioned, our previous analysis assumes implicitly that the monetary amounts accrued to the experimental subjects are perceived by the subjects as to be immediately consumed when they become available. Since however we are only able to observe choice, but unable to observe the subjects’ perceptions or even the consumption itself, it seems reasonable to study the implications of considering the intertemporal budget sets that each

²⁹The four corresponding 95% exact confidence intervals for the odds ratios are [2.42, 167.4], [2.03, 159.1], [2.43, 80.3] and [2.86, 113])

sequence of monetary rewards generates, and discuss the robustness of our previous analysis to these considerations.³⁰


In our experiment, subjects were asked to choose between alternative income streams, all with the same total reward amounts. Any sequence front loaded with the higher payments can be ‘transformed’ into one with larger later payments, so that the implicit experimental interest rate at which the various amounts can be traded off is, in our experiment, zero. On the other hand, considering opportunities outside the lab, experimental subjects could invest all or part of their early monetary receipts to boost future consumption, or use all or part of the future receipts as security towards borrowing to increase current consumption, thus generating an intertemporal budget set. As shown by Cubitt and Read [6] in the context of choices between alternative date-outcome pairs, if the experimental interest rate is either larger or smaller than both the borrowing and the lending rates (with the former assumed higher than the latter), choices might be uninformative: if the budget frontiers corresponding to each date outcome pair are nested, then any standard textbook utility maximiser with convex preferences will always pick the alternative that guarantees the larger budget set.

Extending the argument to the case of our sequences is straightforward, so that with nested budget frontiers in our case, too, observed choice could be completely divorced from time preferences. The budget sets induced by each sequence of rewards for the case of short sequences are depicted in figure 6.

The sets delimited by the dashed lines are those induced by assuming the possibility that subjects can borrow and lend (with borrowing rates higher than the lending rates). In this case, though, we should still observe rational agents with convex preferences over consumption goods to select D over K over I over J in both the case of long and short sequences, while we don’t.³¹ While we agree that this argument might be relevant in the case of sequences of substantial payments, in our case the amounts and time intervals involved are small enough that in practice anticipating that subjects would invest for six months their extra €8 euros (and for a very small return, even brushing aside any transaction costs) seems very far fetched, and the possibility that they might lend this sum at an interest seems even

³⁰We are grateful to a referee for directing us to these issues.

³¹See table 4, where only 55% of the subjects exhibit choice profiles $D \succ K \succ I$ in both short and long sequences. If we consider also the J sequence, then the percentage of agents with the profile $D \succ K \succ I \succ J$ for both sequence lengths falls to 47%, leaving the majority of observed choices unexplained.



wilder.³²


A possibility that seems more difficult to rule out is that this money might be lent or borrowed against in the informal market (e.g. friends). In this case however the money would be exchanged without interest, which tallies with our implicit experimental interest rate. All the sequences generate the same budget set, and lie strung along the (same) budget frontier (the -45 degree line). This also implies that a ‘textbook’ decision maker with strictly convex preferences cannot be indifferent between any two alternatives.

A second and more intriguing possibility is that subjects, unable or unwilling to borrow against future experimental income, also ignore the interest they might earn outside the lab by reinvesting what they earn inside it. Dealing with budget sets allows a decision maker to pick an alternative that would allow him to select the exact consumption bundle that would maximise his utility subject to the budget constraint. For instance, compare I2 and D2, and suppose the agent wants to consume the increasing sequence (20, 28). He must still choose D2 over I2 to achieve this.

To pursue this point further, consider a rational agent, i.e. an agent with convex preference over consumption bundles who is intent in maximising his utility, whose preference for

³²The possibility that the choice of experimental subjects may be driven by their consideration of investment opportunities outside the lab in the context of the elicitation of time preferences was first highlighted by Coller and Williams [5]. In their experiments, however, subjects’ choices involved amounts upwards of \$500 (at 1990s prices).

money derives from wanting to spend it and who can freely hold money. Then sequence D2 must weakly dominate I2, K2 and J2, regardless of the agent's time preferences. The reason is that each of I2, K2 and J2 can be obtained from D2 by just holding money (assumed this to be a costless activity) over the intervening period. For example, by keeping €16 euros a subject with D2 can obtain I2. Whatever the optimal spending pattern for a subject with I2, the very same pattern is also achievable with D2. So, there can be no reason to strictly prefer I2 to D2 (for a rational agent who wants money for its purchasing power and can freely hold it). Could it be that an agent chooses I2 over D2 because he is indifferent between the two options? In this case our experimental subjects would be faced with the budget sets delimited by the solid lines in figure 6.



So the budget set generated by the sequence D contains the one generated by sequence K, and so on. The convexity of preferences over consumption requires any agent that is indifferent between I2 and D2 to be also indifferent between I2 and K2 and between K2 and D2. Thus, conditioning on having chosen I2 over D2, we should expect no particular pattern of choice in the binary comparisons between I2 and D2 and between K2 and D2. Yet, e.g. for the case of the PAY treatment, this is what we observe:

More in general, since we are preventing our experimental subjects from expressing indifference between two or more options, we should be careful in our review of the data. Indeed,

	D2 \succ K2, K2 \succ I2	K2 \succ D2, K2 \succ I2	D2 \succ K2, I2 \succ K2	K2 \succ D2, I2 \succ K2
K2 \succ I2	4	11	0	6

Table 10: Profile of binary choice conditional on having selected I2 over D2.

the analysis of the previous section has shown the various associations between choices. Here however we go back to those results to show them from a slightly different perspective to address the issue of indifference.

Consider the choice between K and D , and suppose the agent is indifferent between these two alternatives, and only selects K because he is forced to a single choice. With convex preferences, there are only two possibilities: either K is also indifferent to I ; or both K and D are strictly preferred to I . If the former, then it must also be that I is indifferent to D , so that we should expect those subjects who chose K over D to more or less distribute evenly in their choices between D and I and between K and I . Yet, inspection of the bottom rows in the two panels in tables 7 reveals a pattern that is not random. Similarly, if both D and K are strictly preferred to I , we should observe a concentration of choices in the bottom-left cell of both tables, which in fact only contains around half of the observations.

Finally, consider the choice between K and I , again assuming that I is chosen over K in spite of the decision maker being indifferent between these two options. A rational agent should then be also indifferent between K and D and between I and D , while what we observe is that most of those agents choosing I over K do not choose D in comparisons involving either I or K .

To summarise, once the assumption that monetary receipts are immediately translated into consumption is dropped, observing that in our experimental setup the implicit interest rate was set equal to zero, the following considerations apply:

- a. if subjects can freely lend and borrow on capital markets, or if, though unable to lend and/or borrow, they choose to store money, then they face a kinked (or truncated) budget constraint. In this case the choices of rational agents do not reveal anything about their time preferences. However rational agents should in this case choose D over K over I over J for both sequence lengths, and less than 50% of our subjects display such behaviour.

- b. if we allow subjects to lend and borrow on the informal market (e.g. from friends) at no interest, then all choices induce the same budget constraint. The predictions of the $(\sigma - \delta)$ model appear consistent with relaxing the assumption that subjects are indifferent between two or more of the available plans.

To conclude, a caveat. In our simple setting assuming a generic preference for increasingness as secondary criterion can be a reduced form for many alternative explanations: for instance, choosing an increasing sequence as a commitment device not to spend it. What we argued is that on its own a preference for increasingness explains very little of the data, and the ‘good news’ is that the modal behaviour conforms to the standard modeling approach. On the other hand, not only are a great deal of data left unexplained within the standard approach, but we have also uncovered patterns of association in choice that are compatible with the $(\sigma - \delta)$ model and incompatible with alternative explanations, thereby reinforcing the general theme that departures from pure discounting cannot be explained by random errors.

7 Concluding remarks

We have proposed a simple model of choice between sequences of monetary rewards with exponential discounting as one of its core elements, the other core element being a ‘secondary heuristics’. This ‘hybrid model’ which combines the traditional, consistent discounting theory with a heuristics component, is very successful at explaining choices between time sequences with an obvious trend (increasing, constant or decreasing). Neither a pure discounting model (of any type) nor a pure heuristics model can explain the data well (though discounting alone does much better than heuristics alone). Of course, our data do not provide *specific* support for the discounting component of the model: other theories, including ones along the lines discussed in section 6, or perhaps hyperbolic discounting itself, might also be used in conjunction with the secondary heuristics. The point we are making is rather that, provided a secondary heuristics is used as we suggest, there is no need to modify the discounting component in order to explain anomalies or observed associations between choices.

Encouragingly, the general pattern of choice we have uncovered in our experiment is consistent with data found elsewhere. Notably, in Gigliotti and Sopher [8], depending on treatment, the choice profiles of up to 80%-90% of their experimental subjects fall into the five choice profiles compatible with our $(\sigma - \delta)$ model.

We conclude with a few comments on the ‘context-dependence’ of the secondary heuristics. Does that mean that our theory is ‘ad hoc’, because we are free to tailor the secondary criterion to the data set we are trying to explain? For example, we can use outcome or time prominence when studying date-outcome pairs, we can use Pareto dominance when studying Rubinstein’s [22] experiment, and we can use sequence trend when studying sequences. However, there is nothing specially ‘ad hoc’ about this. As we have already remarked, at the abstract level, our model departs from the standard choice theoretic model in just one way, by positing two sequential incomplete (but transitive) preference relations instead of a complete one (see Manzini and Mariotti [14] for a general model). So, the secondary criterion is no more context-dependent than any preference relation is: different preference relations will apply (by definition) to different sets of objects. We are not arguing here that different rankings ought to apply to the *same* objects in different contexts (though we do not exclude this possibility). Nonetheless, it is true that - because we interpret the secondary criterion as a heuristic tied to some *salient* feature of the objects - it will generally be easier to glean intuition about an individual’s secondary criterion than about his general preferences. We view this as a strength of the approach, since it makes the abstract model more easily adaptable to specific circumstances.

From our perspective, the search to uncover the nature of the secondary heuristic in cases different from those considered so far ought to be one of the main empirical developments of the theory proposed in this paper.

References

- [1] Abdellaoui, M., A. E. Attema and H. Bleichrodt (2009) ‘Intertemporal tradeoffs for gains and losses: an experimental measurement of discounted utility’, forthcoming on *Economic Journal*.

- [2] Agresti, A., 2002, *Categorical Data Analysis*, Wiley Series in Probability and Statistics, Second Edition, John Wiley and Sons.
- [3] Andersen, S., G. W. Harrison, M. I. Lau and E. E. Rutström, 2008, ‘Eliciting Risk and Time preferences’, *Econometrica*, 76: 583-618.
- [4] Chapman, G. B, 1996, ‘Expectations and preferences for sequences of health and money’, *Organizational Behavior and Human Decision Processes*, vol. 67: 59-75.
- [5] Coller and M. Williams (1999) “Eliciting Individual Discount Rates” *Experimental Economics*, 2: 107-127.
- [6] Cubitt, R. P. and D. Read, 2007, ‘Can intertemporal choice experiments elicit time preferences for consumption?’, *Experimental Economics*, 10: 369–389.
- [7] Frederick, S., G. Loewenstein and T. O’Donoghue, 2002, ‘Time discounting and time preferences: a critical review’, *Journal of Economic Literature*, 40: 351-401.
- [8] Gigliotti, G. and B. Sopher, 1997, ‘Violations of Present Value Maximization in Income Choice’, *Theory and Decision*, 43: 45-69.
- [9] Guyse, J., L. Keller and T. Epple, 2002, ‘Valuing Environmental Outcomes: Preferences for Constant or Improving Sequences’, *Organizational Behavior and Human Decision Processes*, 87: 253-277.
- [10] Loewenstein, G. and D. Prelec, 1991, ‘Negative Time Preference’, *American Economic Review*, 81: 347-352.
- [11] Loewenstein, G. and D. Prelec, 1992, ‘Anomalies in intertemporal choice: evidence and interpretation’, *Quarterly Journal of Economics*, 107: 573-597.
- [12] Loewenstein, G. and D. Prelec, 1993, ‘Preferences for sequences of outcomes’, *Psychological Review*, 100: 91-108.
- [13] Loewenstein, G. and N. Sicherman, 1991, ‘Do workers prefer increasing wage profiles?’, *Journal of Labour Economics*, 9: 67-84.

- [14] Manzini, P. and M. Mariotti, 2007, ‘Sequentially Rationalizable Choice’, *American Economic Review*, 97: 1824-1839.
- [15] Manzini, P. and M. Mariotti, 2006 ‘Two-stage boundedly rational choice procedures: Theory and experimental evidence, IZA Discussion Paper n. 2341, IZA Bonn.
- [16] Manzini, P. and M. Mariotti, 2009, ‘Choice Over Time’, in P. Anand, P. Pattanaik and C. Puppe (eds.) *Oxford Handbook of Rational and Social Choice*, Oxford University Press.
- [17] Manzini, P. and M. Mariotti, 2006, ‘A vague theory of choice over time’, *Advances in Theoretical Economics*, Vol. 6: No. 1, Article 6, <http://www.bepress.com/bejte/advances/vol6/iss1/art6>.
- [18] Noor, J., 2009, ‘Hyperbolic Discounting and the Standard Model: Eliciting Discount Functions’, *Journal of Economic Theory*, 144: 2077-2083.
- [19] Read, D., 2001, ‘Is Time-Discounting Hyperbolic or Subadditive?’, *Journal of Risk and Uncertainty*, 23:5-32.
- [20] Read, D. and M. Powell, 2002, ‘Reasons for sequence preferences’, *Journal of Behavioral Decision Making*, 15: 433-460.
- [21] Rubinstein, A., 2001, ‘A theorist’s view of experiments’, *European Economic Review*, 45: 615-628.
- [22] Rubinstein, A., 2003, ‘Economics and Psychology’? The case of hyperbolic discounting’, *International Economic Review*, 44: 1207-1216.
- [23] Selten, R., 1991, ‘Properties of a measure of predictive success’, *Mathematical Social Sciences*, 21: 153-167.
- [24] Tversky, A., Sattath, S. and Slovic, P., 1988, ‘Contingent weighting in judgement and choice’, *Psychological Review*, 95: 371–384.
- [25] Varey and D. Kahneman, 1990, ‘Experiences extended across time: evaluation of moments and episodes’ *Journal of Behavioral Decision Making*, 5: 169-185.

A Appendix

Instructions

Please note: you are not allowed to communicate with the other participants for the entire duration of the experiment.

The instructions are the same for all you. You are taking part in an experiment to study intertemporal preferences. The project is financed by the ESRC.

Shortly you will see on your screen a series of displays. Each display contains various remuneration plans worth the same total amount of 48 Euros each, staggered in three, six and nine months installments. For every display you will have to select the plan that you prefer, clicking on the button with the letter corresponding to the chosen plan. (HYP: These remuneration plans are purely hypothetical. At the end of the experiment you'll be given a participation fee of 5 €.) (PAY: At the end of the experiment one of the displays will be drawn at random and your remuneration will be made according to the plan you have chosen in that display).

In order to familiarise yourself with the way the plans will be presented on the screen, we shall now give you a completely hypothetical example, based on a total remuneration of 7 Euros.

Plan A

How much When

3 € in one year

1 € in two years

1 € in three years

2 € in four years

Plan B

How much When

1 € in one year

2 € in two years

3 € in three years

1 € in four years

In this example plan A yields 7 € in total in tranches of 3 €, 1 €, 1 € and 2 € in a year, two years, three years and four years from now, respectively, while plan B yields 7 Euros in total in tranches of 1 €, 2 €, 3 € and 1 € in a year, two years, three years and four years from now, respectively.