Journal of Experimental Psychology: Learning, Memory, and Cognition

Tradeoffs Between Sequences: Weighing Accumulated Outcomes Against Outcome-Adjusted Delays

Daniel Read and Marc Scholten

Online First Publication, May 7, 2012. doi: 10.1037/a0028216

Journal of Experimental Psychology:
Learning, Memory, and Cognition

0278-7393/12/$12.00 DOI: 10.1037/a0028216

Tradeoffs Between Sequences: Weighing Accumulated Outcomes Against Outcome-Adjusted Delays

Daniel Read
Warwick Business School, Coventry, United Kingdom

Marc Scholten
ISPA University Institute, Lisbon, Portugal

We extend the recently proposed *tradeoff model* of intertemporal choice (Scholten & Read, 2010) from choices between pairs of single outcomes to pairwise choices involving two-outcome sequences. The core of our proposal is that choices between sequences are made by weighing accumulated outcomes against outcome-adjusted delays. Thus extended, the tradeoff model offers a unified account of recently discovered anomalies in pairwise choices involving two-outcome sequences, including (a) the hidden-zero effect, in which explicit reference to the zero outcomes of the options increases patience, (b) the front-end amount effect, in which the addition of a front-end amount to both options decreases patience, and (c) the mere token effect, in which the addition of an early outcome to both options increases patience. Not only does the extended tradeoff model accommodate these anomalies, it also correctly predicts (d) violations of independence, (e) a reversal of the front-end amount effect, (f) the effect of relocating the front-end amount to the back end of both options, and (g) a dependence of the "mere" token effect on the magnitude of the token. In quantitative analyses, the extended tradeoff model offers an accurate account of the data.

Keywords: intertemporal choice, discounting, tradeoffs, sequences

Compare the following choices, each between a pair of single outcomes: one smaller but sooner (*SS*) and the other larger but later (*LL*):

Pair A: *SS* $10 today
LL $30 in 2 months

Pair A_C: *SS* $1,010 today
LL $1,030 in 2 months

Pair A_C is constructed from Pair A by adding a constant (C = $1,000) to both outcomes ($10 and $30). The *relative magnitude effect* (Scholten & Read, 2010) is that this decreases the preference for *LL* over *SS*. A simple psychological model in which decision makers compare the interest rate they will earn by choosing *LL* with the minimum interest they want to earn for waiting (which is sometimes called the *pure rate of time preference*) naturally predicts this effect, because adding a common amount to both outcomes reduces the interest rate. The first choice offers a monthly interest rate of about 73%, while the second offers about 1%.[1] On closer scrutiny, however, the interest-rate model fails. Three recently reported results show anomalies to this model when *SS* and *LL* are expanded into elementary sequences of two outcomes.

One sequence effect is the *hidden-zero effect* (Magen, Dweck, & Gross, 2008), which is that by making it explicit that *SS* will pay zero when *LL* pays out and that *LL* will pay zero when *SS* pays out, thus turning both options into two-outcome sequences, the preference for *LL* increases. To illustrate, we compare the previously described Pair A with Pair A_Z, in which the zero outcomes are made explicit:

Pair A: *SS* $10 today
LL $30 in 2 months

Pair A_Z: *SS* $10 today and $0 in 2 months
LL $0 today and $30 in 2 months

Choice of *LL* is more likely in Pair A_Z than in Pair A. This is an anomaly to the interest-rate model, because making the zero outcomes explicit leaves the interest rate unchanged.

Another sequence effect is what we call the *front-end amount effect* (Rao & Li, 2011), which is that by adding a common amount to the immediate outcomes of both options, thus turning *LL* into a two-outcome sequence, the preference for *LL* decreases. To illustrate, we compare the now-familiar Pair A with Pair A_F, in which a common amount (F = $1,000) is added to the immediate outcomes of both options:

Daniel Read, Warwick Business School, Coventry, United Kingdom; Marc Scholten, ISPA University Institute, Lisbon, Portugal.

We acknowledge the financial support from the Fundação para a Ciência e Tecnologia (FCT), Program POCI 2010, and Project PTDC/PSI-PCO/101447/2008. We are indebted to João Nogueira.

Correspondence concerning this article should be addressed to Marc Scholten, ISPA University Institute, Rua Jardim do Tabaco 34, Lisboa 1149-041, Portugal. E-mail: scholten@ispa.pt

[1] This "psychological" model is often described as the *normative* model from economics, but the two models differ in that the economic model holds that people will choose to maximize the opportunity cost of money, which entails choosing *LL* if the interest rate offered exceeds the (risk-adjusted) interest rate they can earn from other uses of money. See Read, Frederick, and Scholten (2012) for further details. This economic model predicts that a person will choose on the basis of a single interest rate, just as does the psychological model.

Pair A: *SS* $10 today
LL $30 in 2 months

Pair A_F: *SS* $1,010 today
LL $1,000 and $30 in 2 months

Choice of *LL* is less likely in Pair A_F than in Pair A. This is again an anomaly to the interest-rate model, because adding the front-end amount to both options leaves the interest rate unchanged: The interest rate derives from the ratio between how much more *LL* offers later and how much more *SS* offers sooner, and this ratio is $30/$10 in both choices.[2] Note that, in choices involving sequences, *SS* is the option that yields more in the first period but less across the two periods.

A third sequence effect is the *mere token effect* (Urminsky & Kivetz, 2011), which is that by adding a sooner outcome (the token) to both options, thus turning both options into two-outcome sequences, preference for *LL* increases. To illustrate, compare the following option pairs:

Pair B: *SS* $300 in 1 week
LL $900 in 1 year

Pair B_T: *SS* $50 in 3 days and $300 in 1 week
LL $50 in 3 days and $900 in 1 year

Choice of *LL* is more likely in Pair B_T than in Pair B. This is once more an anomaly to the interest-rate model, because the common outcome cancels out in the computation of the interest rate.

Our goal in this article is to develop a theoretical account of these three phenomena and to provide new evidence to support that account. We extend the recently proposed *tradeoff model* of intertemporal choice (Scholten & Read, 2010) from choices between pairs of single outcomes to pairwise choices involving elementary sequences of two outcomes. In its original form, the tradeoff model accommodates the full range of anomalies to the interest-rate model in choices between pairs of single outcomes, many of which are also anomalies to established models like the discounted utility model (Samuelson, 1937) and the (quasi-) hyperbolic discounting model (Laibson, 1997; Loewenstein & Prelec, 1992).[3] In this article, we show that the extended tradeoff model accurately accounts for anomalies in choices between a single outcome and a two-outcome sequence and choices between pairs of two-outcome sequences.

The Tradeoff Model

Original Tradeoff Model

Conceptually, the tradeoff model weighs time against outcome, whereas the interest-rate model and other discounting models weigh outcome by time. Consider a choice between a pair of single outcomes, as in Pairs A and A_C. In abstract notation, these are choices between (x_S, t_S) and (x_L, t_L), where $x_L > x_S > 0$ and $t_L > t_S > 0$. In Pair A, x_S and x_L are $10 and $30, and t_S and t_L are "today" and "in 2 months." As described by the tradeoff model, *SS* has an advantage over *LL* along the time attribute (the difference between delays: It pays 2 months sooner), whereas *LL* has an advantage over *SS* along the outcome attribute (the difference between outcomes: It pays $20 more). These advantages are not differences between *raw* attribute amounts; rather, they are differences between *weighted* delays, $w(t_L) - w(t_S)$, and differences between *valued* outcomes, $v(x_L) - v(x_S)$. The decision maker prefers *SS* when the time advantage is greater, prefers *LL* when the outcome advantage is greater, and is indifferent between the options when the time advantage equals the outcome advantage:

$$v(x_L) - v(x_S) = \kappa[w(t_L) - w(t_S)], \quad (1)$$

where $\kappa > 0$ is a tradeoff parameter, which scales the difference between weighted delays and the difference between valued outcomes to a common currency.[4]

The value function v and the time-weighing function w are reference-dependent functions ranging from identity functions, that is, $v(x) = x$ and $w(t) = t$ (constant sensitivity) to zero functions, that is, $v(x) = 0$ for all x, and $w(t) = 0$ for all t (insensitivity). Between these two limits, v and w are concave functions, thus exhibiting diminishing (absolute) sensitivity (see Scholten & Read, 2010; Tversky & Kahneman, 1991): The marginal impact of an outcome decreases with its magnitude (e.g., adding $10 to $20 has a bigger impact than adding $10 to $200), and the marginal impact of a delay decreases with its length (e.g., adding 1 week to 1 day has a bigger impact than adding it to 1 year).

Extended Tradeoff Model

We extend the tradeoff model from choices between pairs of single outcomes to pairwise choices involving two-outcome sequences and thus accommodate the anomalies described in the introduction and in the experiments that follow. For two-outcome sequences, the options are $SS = (x_{S_1}, t_{S_1}; x_{S_2}, t_{S_2})$ and $LL = (x_{L_1}, t_{L_1}; x_{L_2}, t_{L_2})$, where *SS* offers a greater gain than *LL* in Period 1 (i.e., $x_{S_1} > x_{L_1}$), but *LL* offers a greater gain across the two periods (i.e., $x_{L_1} + x_{L_2} > x_{S_1} + x_{S_2}$. To illustrate, in the choice between "$10 today and $30 in 2 months" and "$5 today and $35 in 2 months," the first sequence is *SS*, because it offers more in Period 1, and the second sequence is *LL*, because it offers more across the two periods.

The key intuition underlying the extended tradeoff model is that people treat a two-outcome sequence as a single dated outcome $(x, \hat{t})$, where x is the total amount offered by the sequence (i.e., $x = x_1 + x_2$), and $\hat{t}$ is the "average" delay, where the averaging process depends on the magnitude of the outcomes. Thus, for instance, the sequence "$500 today and $30 in 2 months" might be treated as the single dated outcome "$530 in 2 weeks." The 2-week delay of

[2] For the choice between "$1,010 today" and "$1,000 today and $30 in 2 months," the interest rate would be computed as $r = \left(\frac{\$30 - \$0}{\$1{,}010 - \$1{,}000}\right)^{1/(2-0)} - 1$.

[3] For documentation and discussion of these anomalies, see Leland (2002), Read (2001), Roelofsma and Read (2000), Rubinstein (2003), and Scholten and Read (2006, 2010), among others.

[4] In the original, and more general, statement of the tradeoff model, κ is a parameter of a nonlinear tradeoff function (Scholten & Read, 2010, 2012a, 2012b). In Equation 1, this nonlinear tradeoff function is reduced to a simple multiplication by a tradeoff parameter κ. The simplified statement suffices for the current analysis because, by the design of our experiments, in which delays are held constant (at 0 and 2 months), we control for the phenomena accommodated by the nonlinear tradeoff function.

the single dated outcome is an average of the zero delay and the 2-month delay of the sequence, where a greater weight is assigned to the zero delay than to the 2-month delay, because $500 is a larger outcome than $30. The extended tradeoff model is described more precisely, and more exhaustively, by the following propositions:

1. Outcome accumulation. The outcomes of a sequence are first summed and then valued, that is, $v(x_1 + x_2)$.

2. Delay adjustment. The *adjusted delay* to the accumulated outcome is an average of the delays to the constituent outcomes. The average lies between the weighted and unweighted average of the delays, as implied by the following propositions.

2.1. Toward weighted averaging: Outcome-dependent weighing of delays to nonzero outcomes. In weighing delays to nonzero outcomes against one another, the weight of a delay to an outcome increases with the magnitude of the outcome. For instance, the adjusted delay of "$500 today and $530 in 2 months" is longer than the adjusted delay of "$500 today and $30 in 2 months."

2.2. Toward unweighted averaging: Outcome-independent weighing of delays to stated zero outcomes. In weighing a delay to a stated zero outcome against a delay to a nonzero outcome, the weight of the delay to the stated zero outcome is independent of the magnitude of the nonzero outcome. For instance, the adjusted delay of "$0 today and $30 in 2 months" is shorter than the (unadjusted) delay of "$30 in 2 months," but no shorter than the adjusted delay of "$0 today and $530 in 2 months."

Formally, Propositions 2.1 and 2.2 imply a single averaging rule. According to this rule, the adjusted delay is

$$\hat{t} = \frac{(qx_2 + x_1)t_1 + (qx_1 + x_2)t_2}{(1 + q)(x_1 + x_2)} \quad \text{if} \quad x_1, x_2 \geq 0, \tag{2}$$

where q ranges from 0 ($\hat{t}$ is the weighted average of t_1 and t_2) to 1 ($\hat{t}$ is the unweighted average of t_1 and t_2). When x_1 is a stated zero outcome, Equation 2 reduces to

$$\hat{t} = \frac{qt_1 + t_2}{1 + q},$$

which is independent of amount, and lies between t_1 and t_2. Conversely, when x_2 is a stated zero outcome, Equation 2 reduces to

$$\hat{t} = \frac{t_1 + qt_2}{1 + q}.$$

When x_1 is an unstated zero outcome, $\hat{t} = t_2$; conversely, when x_2 is an unstated zero outcome, $\hat{t} = t_1$.

In addition to these propositions, which are original with the extended tradeoff model, we introduce a third that is already well established. This third proposition describes a preference pattern that emerges in our experiments as well.

3. Preference for spreading (Loewenstein & Prelec, 1993). Deviation from a uniform distribution of outcomes detracts from the value of the accumulated outcome; that is, $v(x_1 + x_2) - \sigma d(x_1, x_2)$, where d is the deviation from a uniform distribution, or, following Loewenstein and Prelec (1993), half the absolute deviation between the outcomes, that is, $\frac{1}{2}|x_1 - x_2|$, and $\sigma > 0$ is preference for spreading. For instance, in the choice between "$20 today and $0 in 2 months" and "$10 today and $30 in 2 months," *SS* and *LL* deviate equally from a uniform distribution ($d = 10$). Alternatively, in the choice between "$10 today and $10 in 2 months" and "$0 today and $40 in 2 months," *LL* deviates from a uniform distribution ($d = 20$), but *SS* does not ($d = 0$).

Tradeoff rule. Given these propositions, the decision maker will, in the extended tradeoff model, be indifferent between *SS* and *LL* when

$$v(x_{L_1} + x_{L_2}) - v(x_{S_1} + x_{S_2}) - \sigma(d(x_{L_1}, x_{L_2}) - d(x_{S_1}, x_{S_2})) = \kappa(w(\hat{t}_L) - w(\hat{t}_S)), \tag{3}$$

where the adjusted delays, $\hat{t}_L$ and $\hat{t}_S$, are given by Equation 2. When both x_{S_2} and x_{L_1} are unstated zero outcomes, the extended tradeoff model in Equation 3 reduces to the original tradeoff model in Equation 1. We next conducted a series four experiments in which we tested implications of the extended tradeoff model.

Experiment 1: Violation of Independence

In Experiment 1, we examined a manipulation that pits the effect of delay adjustment against the effect of diminishing sensitivity to accumulated outcomes. The net result of this manipulation is a *violation of independence*, analogous to the ones obtained by Loewenstein (1987) and Loewenstein and Prelec (1993). This violation of independence, however, favors the tradeoff model over alternative models.

Given a choice between two single outcomes, (x_S, t_S) and (x_L, t_L), we construct a choice between two sequences by inserting, in both options, an intermediate outcome, x_M, available after an intermediate delay t_M. Conventional discounting models predict that the common consequence will not affect choice, because it cancels out in the comparison between *SS* and *LL*. An additional aspect of our manipulation is that x_M lies exactly between x_S and x_L; that is., $x_L - x_M = x_M - x_S$. This neutralizes the preference for spreading, because $d(x_M, x_L) = d(x_S, x_M)$, and any preference for improvement, as identified by Loewenstein and Prelec's (1993) model of preferences over sequences. Thus, current models predict that the common consequence will not affect choice.

The tradeoff model predicts that the common consequence can either decrease or increase the preference for *LL*, depending on the relative contribution of two processes. On the one hand, the sensitivity to the outcome difference between *LL* and *SS* diminishes, that is, $v(x_L + x_M) - v(x_S + x_M) < v(x_L) - v(x_S)$, which decreases the preference for *LL*. On the other hand, the delay of *SS* increases (i.e., $\hat{t}_S > t_S$), whereas the delay of *LL* decreases (i.e., $\hat{t}_L < t_L$), which increases the preference for *LL*. If one process outweighs the other, there will be a violation of independence.

In Experiment 1, the outcomes are $x_S = \$300$, $x_M = \$350$, and $x_L = \$400$, and the delays (in weeks) are $t_S = 0$, $t_M = 4$, and $t_L = 50$. With these outcomes and delays, the effect of diminishing sensitivity to accumulated outcomes is small relative to the effect of delay adjustment. Figure 1 shows how sensitivity to the outcome difference between *LL* and *SS* diminishes with the addition of x_M to both x_S and x_L. Diminishing sensitivity falls between two limits: Constant sensitivity (the valued outcome difference is $x_L - x_S = \$100$, regardless of x_M) and insensitivity (the valued outcome difference is $0, regardless of x_M). Between these limits, sensitivity to the outcome difference between *LL* and *SS* diminishes with

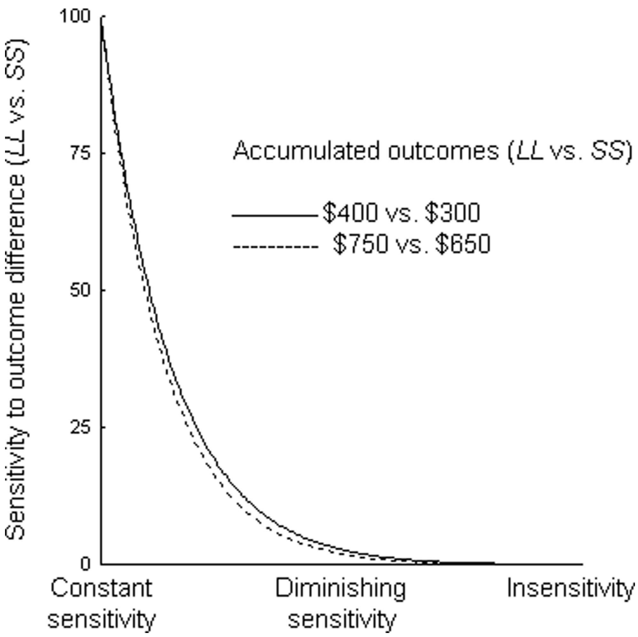

Figure 1. Experiment 1: Sensitivity to outcome differences as a function of the accumulated outcomes of *SS* (smaller but sooner) and *LL* (larger but later).

the addition of x_M, but it can be seen that on the scale from $0 to $100, it diminishes very little. On the other hand, the effect of delay adjustment is substantial. The adjusted delay of *SS* will lie between $t_S = 0$ and $t_M = 4$, and the adjusted delay of *LL* will lie between $t_M = 4$ and $t_L = 50$. With the addition of x_M at t_M, the delay of *SS* increases from $t_S = 0$ to $\hat{t}_S \approx 2$, whereas the delay of *LL* decreases from $t_L = 50$ to $\hat{t}_L \approx 28$.[5] Therefore, the adjusted delays are a lot closer than the unadjusted ones. We thus predicted that the addition of x_M at t_M would lead to an increased preference for *LL*.

Method

A total of 132 workers on Amazon Mechanical Turk (U.S. residents, 39% male, average age of 46 years, 98% having at least attended college or university, and 59% being employed) participated by completing an online questionnaire related to several studies. The present study included two items, one involving single outcomes and the other involving two-outcome sequences. Each participant responded to both items, in an order randomized across participants.

Results

The results for the single outcomes were as follows:

SS: Receive $300 today. [80%]
LL: Receive $400 in 50 weeks. [20%]

The results for the two-outcome sequences were as follows:

SS: Receive $300 today and receive $350 in 4 weeks. [73%]
LL: Receive $350 in 4 weeks and receive $400 in 50 weeks. [27%]

Thus, in violation of independence, the common consequence of $350 in 4 weeks increased the preference for *LL*, $\chi^2(1) = 3.86$, $p < .005$. According to the tradeoff model, this result shows that delay adjustment outweighed diminishing sensitivity to accumulated outcomes.

Experiment 2: Zero Outcomes and Front-End Amounts

While Experiment 1 pitted the effect of delay adjustment against the effect of diminishing sensitivity to accumulated outcomes, Experiment 2 introduced preference for spreading and applied all three elements of the tradeoff model to the hidden-zero effect and the front-end amount effect.

Hidden-Zero Effect

The hidden-zero effect is that making zero outcomes explicit increases preference for *LL*. We examine two scenarios. One is the introduction of explicit-zero outcomes to a choice between two single outcomes, as in Magen et al. (2008). Consider Pair A in Table 1. The choice between "$10 today" and "$30 in 2 months," in which both *SS* and *LL* are single outcomes, becomes a choice between "$10 today and $0 in 2 months" and "$0 today and $30 in 2 months," in which both *SS* and *LL* are sequences. The hidden-zero effect is, in this scenario, the net result of two processes. On the one hand, the explicit-zero outcomes introduce a situation in which *LL* deviates more from a uniform distribution than *SS* ($^1/_2|0 - 30| = 15 > {}^1/_2|10 - 0| = 5$), and this decreases the preference for *LL*. On the other hand, the explicit-zero outcome increases the delay of *SS* (for $q = .3$, from 0 to $\approx {}^1/_2$) and decreases the delay of *LL* (from 2 to ≈ 1.5), which increases the preference for *LL*. The hidden-zero effect occurs when delay adjustment outweighs preference for spreading.

The other scenario is the introduction of an explicit-zero outcome to a choice between a single outcome and a sequence. Consider Pair B in Table 1. The choice between "$510 today" and "$500 today and $30 in 2 months," in which *SS* is a single outcome and *LL* is a sequence, becomes a choice between "$510 today and $0 in 2 months" and "$500 today and $30 in 2 months," in which both *SS* and *LL* are sequences. The hidden-zero effect is, in this scenario, the *joint* result of two processes: The explicit-zero outcome introduces a situation in which *SS* deviates from a uniform distribution ($d = 255$), and increases the delay of *SS* (for $q = .3$, from 0 to ≈ 0.5). Both of these processes increase the preference for *LL*, thus producing the hidden-zero effect.

Front-End Amount Effect

The front-end amount effect, as reported by Rao and Li (2011), is that adding a common amount to the immediate outcomes of

[5] Because x_M is very similar to x_S and x_L, the weighted averages of t_S and t_M, and of t_M and t_L will be very similar to the unweighted averages. The adjusted delay of *SS* is ≈ 2, regardless of the averaging rule. The adjusted delay of *LL* lies between 27 (unweighted average) and ≈ 28.5 (weighted average). If in Equation 2, $q = .3$, which is a value close to the estimates that we obtain later on in a quantitative analysis, the adjusted delay of *LL* is ≈ 28 weeks.

Table 1
Adding a Common Amount ($A_{0;2}$) to the Immediate Outcome of Smaller But Sooner (SS) and the Delayed Outcome of Larger but Longer (LL) and Adding a Common Amount (A_0) to the Immediate Outcomes of Both Options in the Explicit-Zero Condition

	$A_{0;2}$		
A_0	€0	€500	€1,000
€0	**A** (€10, 0; €0, 2) (€0, 0; €30, 2) 73.21%	**D** (€510, 0; €0, 2) (€0, 0; €530, 2) 1.94%	**G** (€1,010, 0; €0, 2) (€0, 0; €1,030, 2) 0.99%
€500	**B** (€510, 0; €0, 2) (€500, 0; €30, 2) 73.21%	**E** (€1,010, 0; €0, 2) (€500, 0; €530, 2) 1.94%	**H** (€1,510, 0; €0, 2) (€500, 0; €1,030, 2) 0.99%
€1,000	**C** (€1,010, 0; €0, 2) (€1,000, 0; €30, 2) 73.21%	**F** (€1,510, 0; €0, 2) (€1,000, 0; €530, 2) 1.94%	**I** (€2,010, 0; €0, 2) (€1,000, 0; €1,030, 2) 0.99%

Note. Implicit-zero condition is obtained by suppressing zero outcomes. Delays are in months. Percentages are monthly interest rates.

both options decreases the preference for *LL*. In Table 1, this happens when moving from Pair A to Pair B. In the tradeoff model, the front-end amount effect is the net result of three processes. Two of these, delay adjustment and diminishing sensitivity, apply equally to the explicit- and implicit-zero conditions. The third, preference for spreading, operates differently in these two conditions.

Delay adjustment and diminishing sensitivity. On the one hand, the front-end amount does not change the delay of *SS*, but it decreases the delay of *LL*, which increases the preference for *LL*.[6] On the other hand, with the introduction of the front-end amount, the sensitivity to the outcome difference between *LL* and *SS* diminishes, which decreases the preference for *LL*.[7]

Preference for spreading. In the implicit-zero condition, the front-end amount changes a situation in which both options are single outcomes, and therefore are not adversely affected by the deviation from a uniform distribution, to a situation in which *SS* remains a single outcome but *LL* becomes a sequence, and thus becomes affected by the deviation from a uniform distribution ($d = 235$). This decreases the preference for *LL*. In the explicit-zero condition, the front-end amount changes a situation in which *LL* deviates more from a uniform distribution than *SS* ($15 > 5$) into one in which the reverse is true ($235 < 255$). This increases the preference for *LL*.

Conclusion. In the implicit-zero condition, the front-end amount effect occurs when diminishing sensitivity and preference for spreading outweigh delay adjustment; in the explicit-zero condition, it occurs when diminishing sensitivity outweighs preference for spreading and delay adjustment.

Reversal of the Front-End Amount Effect

The tradeoff model also points to the possibility that the front-end amount effect may *reverse*. We next compared Pairs A–B, in which, drawing on the results reported by Rao and Li (2011), we expected the front-end amount effect, with Pairs D–E, in which the front-end amount effect may reverse.

Implicit-zero condition. In this condition, the front-end amount effect reverses when delay adjustment outweighs diminishing sensitivity and preference for spreading. When moving from Pairs A–B to Pairs D–E, the effects of all three processes are attenuated, and any reversal of the front-end amount effect depends on whether delay adjustment outweighs diminishing sensitivity and preference for spreading.

By diminishing sensitivity, the front-end amount decreases preference for *LL*. However, sensitivity diminishes at a diminishing rate: Adding €500 to €30 and €10 (in Pairs A–B) leads to a greater decrease in the sensitivity to the outcome difference than adding another €500 to €530 and €510 (in Pairs D–E), as shown in Figure 2. Therefore, when moving from Pairs A–B to Pairs D–E, the negative effect of diminishing sensitivity on the preference for *LL* is attenuated.

By preference for spreading, the front-end amount decreases preference for *LL*. In each option pair, *SS* is a single outcome, and therefore is not adversely affected by the deviation from a uniform distribution. In Pairs A and D, *LL* is also a single outcome, and is therefore not affected either by the deviation from a uniform distribution. In Pairs B and E, however, *LL* is a sequence, and one that deviates more from a uniform distribution in Pair B ($d = 235$)

[6] In the implicit-zero condition, *SS* is a single outcome, so that there is no delay adjustment. In the explicit-zero condition, *SS* is a sequence of a positive outcome and a zero outcome, so that by the proposition of *outcome-independent weighing of delays to stated zero outcomes*, the adjusted delay of *SS* is unaffected by the magnitude of the positive outcome.

[7] The front-end amounts employed by Rao and Li (2011) ranged from huge (hundreds of thousands of yuans) to gigantic (hundreds of billions of yuans). As discussed next, it is not necessary to employ such numbers in order to obtain the front-end amount effect.

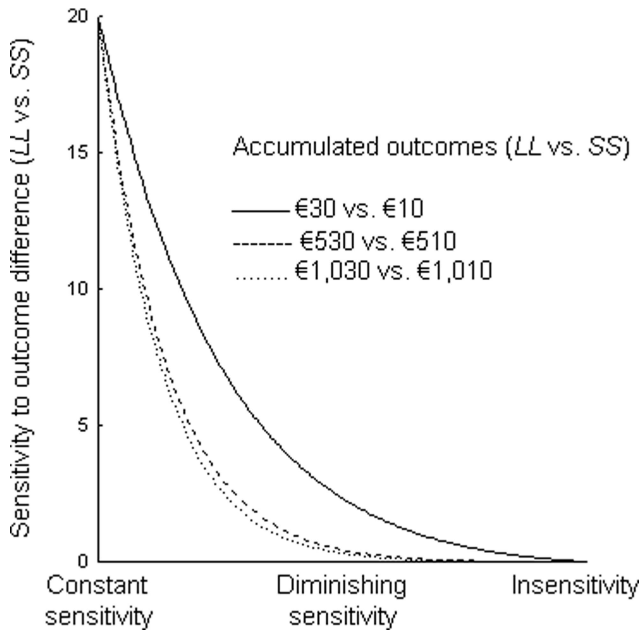

Figure 2. Experiments 2 and 3: Sensitivity to outcome differences as a function of the accumulated outcomes of *SS* (smaller but sooner) and *LL* (larger but later).

than in Pair E (d = 15). Thus, when moving from Pairs A–B to Pairs D–E, the negative effect of preference for spreading on the preference for *LL* is attenuated.

By delay adjustment, the front-end amount increases preference for *LL*. In each option pair, *SS* is a single outcome, so that its delay is not adjusted (0). In Pairs A and D, *LL* is also a single outcome, so that its delay is not adjusted either (2). In Pairs B and E, however, *LL* is a sequence, and one that has a shorter adjusted delay in Pair B (for q = .3, ≈ 0.5) than in Pair E (about 1). Therefore, when moving from Pairs A–B to Pairs D–E, the positive effect of delay adjustment on the preference for *LL* is attenuated.

In sum, when moving from Pairs A–B to Pairs D–E, both the positive effect of delay adjustment on the preference for *LL* and the negative effects of diminishing sensitivity and preference for spreading on the preference for *LL* are attenuated. The front-end amount effect reverses when the attenuated effect of delay adjustment is greater than the attenuated effects of diminishing sensitivity and preference for spreading.

Explicit-zero condition. In this condition, the front-end amount effect reverses when preference for spreading and delay adjustment outweigh diminishing sensitivity. When moving from Pairs A–B to Pairs D–E, the negative effect of diminishing sensitivity on the preference for *LL* and the positive effect of delay adjustment on the preference for *LL* are, as in the implicit-zero condition, attenuated, but the positive effect of preference for spreading on the preference for *LL* is *accentuated*. On the one hand, the difference between *LL* and *SS* in the deviation from a uniform distribution is 15 – 5 = 10 (in favor of *SS*) in Pair A and 235 – 255 = −20 (in favor of *LL*) in Pair B. On the other hand, the difference between *LL* and *SS* in the deviation from a uniform distribution is 265 – 255 = 10 (in favor of *SS*) in Pair D, and 15 − 505 = −490 (in favor of *LL*) in Pair E. Therefore, when moving from Pairs A–B to Pairs D–E, the positive effect of preference for spreading on the preference for *LL* is accentuated. The front-end amount effect reverses when the accentuated effect of preference for spreading and the attenuated effect of delay adjustment are greater than the attenuated effect of diminishing sensitivity.

Method

A total of 277 Portuguese residents (42% male, average age 30 years, 65% having at least completed college or university, and 74% being employed or a student) participated by completing an online questionnaire. Participants were randomly assigned to the implicit-zero condition or the explicit-zero condition. The order of the stimuli in Table 1 was randomized across participants.

Results

Figure 3 shows the choice probabilities for the nine cells of the within-participant design separately for the implicit-zero and the explicit-zero conditions. We conducted a 3 (common amount added to the immediate outcome of *SS* and the delayed outcome of *LL*, denoted as $A_{0;2}$) × 3 (common amount added to the immediate outcomes of both options, denoted as A_0) × 2 (implicit- or explicit-zero outcomes) mixed analysis of variance. Three results emerged.

First, preference for *LL* decreased as $A_{0;2}$ increased, $F(2, 550) = 125.46$, $p < .005$, $\eta_p^2 = .31$. This is the relative magnitude effect, and can be accounted by the interest-rate model, because interest rates decrease as $A_{0;2}$ increases, and by the tradeoff model, because the sensitivity to accumulated outcomes decreases as $A_{0;2}$ increases.

Furthermore, preference for *LL* increased when zero outcomes were stated explicitly, $F(1, 275) = 41.07$, $p < .005$, $\eta_p^2 = .13$. This is the hidden-zero effect and is a replication of the result reported by Magen et al. (2008).

Finally, $A_{0;2}$ interacted with A_0, $F(4, 1100) = 48.76$, $p < .005$, $\eta_p^2 = .15$. For $A_{0;2}$ = €0 (Cells A, B, and C), preference for *LL* decreased as A_0 increased, $F(2, 550) = 17.79$, $p < .005$, $\eta_p^2 = .06$. This is the front-end amount effect and is a replication of the result reported by Rao and Li (2011). For $A_{0;2}$ = €500 (Cells D, E, and F) and €1,000 (Cells G, H, and I), however, preference for *LL* increased as A_0 increased, $F(2, 550) = 61.12$, $p < .005$, $\eta_p^2 = .18$. This is a reversal of the front-end amount effect. Overall, the implications of the tradeoff model were confirmed.[8]

Experiment 3: Front-End and Back-End Amounts

Experiment 2 investigated the net results of delay adjustment, diminishing sensitivity to accumulated outcomes, and preference for spreading; in contrast, Experiment 3 isolated the effect of preference for spreading by comparing the front-end amount condition, in which a common amount is added to the immediate outcomes of both options (see Table 1), with a back-end amount condition, in which the common amount is added to the *delayed*

[8] Other significant results were a main effect of $A_{0;2}$, $F(2, 550) = 19.58$, $p < .005$, $\eta_p^2 = .07$, which was qualified by its interaction with A_0 and a weak and subtle interaction effect between $A_{0;2}$ and the implicit-zero versus explicit-zero condition, $F(2, 550) = 3.17$, $p < .005$, $\eta_p^2 = .01$.

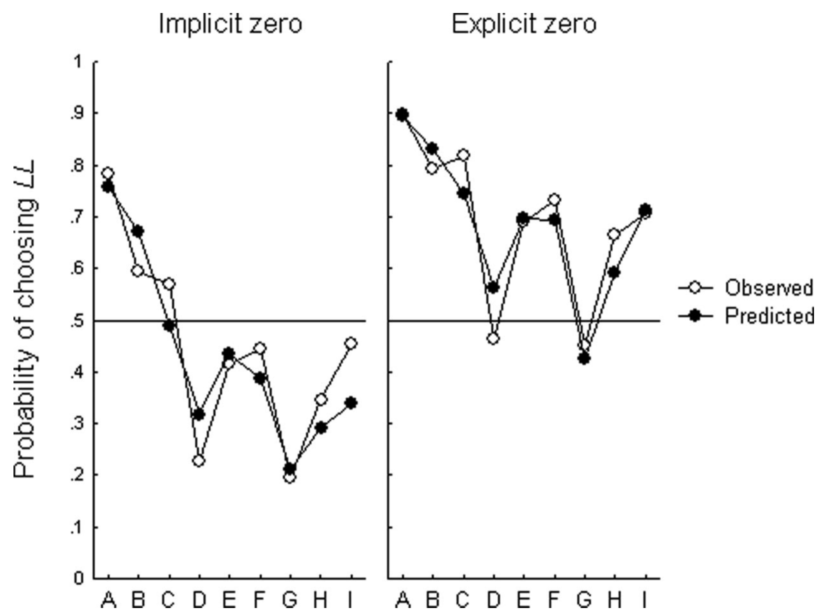

Figure 3. Experiment 2: Observed and predicted probability of choosing *LL* (larger but later) in the nine cells of the within-participant design (from A to I), separately for the implicit-zero and explicit-zero conditions. The width of the confidence intervals ranged from .086 to .117.

outcomes of both options (see Table 2). Pairs A, D, and G are the same in the two conditions, because the common amount is €0. We therefore focus on the other pairs, in which the common amount is either €500 or €1,000.

A comparison between the two tables controls for the effects of diminishing sensitivity and delay adjustment and thus isolates the effect of preference for spreading, as discussed next.

Diminishing Sensitivity

For each choice in Tables 1 and 2, the accumulated outcomes are the same. For instance, in Pair B, the accumulated outcomes are €510 for *SS* and €530 for *LL*, both in the front-end amount condition and in the back-end amount condition. Thus, the impact of diminishing sensitivity is removed.

Delay Adjustment

Relocating the common amount of money from the front end to the back end increases the adjusted delays of *SS* and *LL* by approximately the same amount of time. The greatest difference would be observed for Pair B, and when the adjusted delays are fully weighted averages of the constituent delays: If, in Equation 2, $q = 0$, relocating the common amount from the front end to the back end increases the adjusted delay of *SS* by 1.96 months, and the adjusted delay of *LL* by 1.89 months.[9] If, however, $q = .3$, a value close to the estimates obtained in the next section, the adjusted delay of *SS* increases by 1.06 months and the adjusted delay of *LL* by 1.02 months. Therefore, the impact of delay adjustment is also removed.

Preference for Spreading

With the effects of diminishing sensitivity and delay adjustment controlled for, the comparison between the front-end and back-end amount conditions isolates the effect of preference for spreading: In the front-end amount condition, *SS* deviates more from a uniform distribution than *LL* (for Pair B, 255 > 235), but, in the back-end amount condition, the reverse is true (for Pair B, 265 > 245). Therefore, choice of *LL* should be less likely in the back-end amount condition than in the front-end amount condition.

Method

A total of 470 Portuguese residents (38% male, average age 40 years, 81% having at least completed college or university, and 85% being employed or a student) participated by completing an online questionnaire.[10] Participants were randomly assigned to the front-end amount condition and the back-end amount condition. The order of the stimuli in Tables 1 and 2 was randomized across participants.

Results

Figure 4 shows the choice probabilities for the nine cells of the within-participant design, separately for the front-end amount and back-end amount conditions. We conducted a 2 (front-end amount or back-end amount condition) × 3 (common amount added to the immediate outcomes of both options, denoted as A_0, or the delayed outcomes of both options, denoted as A_2) × 3 (common amount added to the immediate outcome of *SS* and the delayed outcome of

[9] The increase in the adjusted delay is $\Delta = \frac{2(1-q)A}{(1+q)(A+x_1+x_2)}$, where A is the amount being relocated, x_1 and x_2 are the other amounts in the sequence, and q is the departure from weighted averaging.

[10] The results of this experiment were replicated with 276 U.S. residents working on Amazon Mechanical Turk. The additive constants used in this replication were $0, $80, and $480.

Table 2
Adding a Common Amount ($A_{0;2}$) to the Immediate Outcome of Smaller but Sooner (SS) and the Delayed Outcome of Larger but Longer (LL) and Adding a Common Amount (A_2) to the Delayed Outcomes of Both Options

	$A_{0;2}$		
A_2	€0	€500	€1,000
€0	**A** (€10, 0; €0, 2) (€0, 0; €30, 2) 73.21%	**D** (€510, 0; €0, 2) (€0, 0; €530, 2) 1.94%	**G** (€1,010, 0; €0, 2) (€0, 0; €1,030, 2) 0.99%
€500	**B** (€10, 0; €500, 2) (€0, 0; €530, 2) 73.21%	**E** (€510, 0; €500, 2) (€0, 0; €1,030, 2) 1.94%	**H** (€1,010, 0; €500, 2) (€0, 0; €1,530, 2) 0.99%
€1,000	**C** (€10, 0; €1,000, 2) (€0, 0; €1,030, 2) 73.21%	**F** (€510, 0; €1,000, 2) (€0, 0; €1,530, 2) 1.94%	**I** (€1,010, 0; €1,000, 2) (€0, 0; €2,030, 2) 0.99%

Note. Delays are in months. Percentages are monthly interest rates.

LL, denoted as $A_{0;2}$) mixed analysis of variance. Two results emerged.

First, preference for *LL* decreased when $A_{0;2}$ increased, $F(2, 936) = 237.16$, $p < .005$, $\eta_p^2 = .34$. This is the relative magnitude effect, which can be accounted for by the interest-rate model and the tradeoff model.

Second, choice of *LL* was less likely in the back-end amount condition than in the front-end amount condition, $F(1, 468) = 21.61$, $p < .005$, $\eta_p^2 = .04$. This main effect was qualified by an interaction effect between condition and the common amount added in each condition, $F(2, 936) = 65.71$, $p < .005$, $\eta_p^2 = .12$: Choice of *LL* was less likely in the back-end amount condition than in the front-end amount condition for A_0, A_2 = €500 (Cells B, E, and H) and €1,000 (Cells C, F, and I), that is, when the pairs differed between the conditions, $F(1, 468) = 48.05$, $p < .005$, $\eta_p^2 = .09$, but not for A_0, A_2 = €0 (Cells A, D, and G), that is, when the pairs did not differ between the conditions, $F(1, 468) = 2.17$, $p > .10$, $\eta_p^2 = .00$. The difference between the conditions confirms the effect of preference for spreading. Thus, the implications of the tradeoff model were again confirmed.[11]

Quantitative Analysis

So far, we have shown that the tradeoff model can offer a qualitative account of the data. Because Experiments 2 and 3 each provide 18 data points, it is feasible to examine whether it can also offer a quantitative account of the data. We estimated the tradeoff model on each data set, using the full specification in Equation 3. Details about the estimation of the tradeoff model are given in the Appendix. Parameter estimates and goodness of fit are reported in Table 3. The parameter estimates are similar across experiments (with $q \approx .3$), and the proportion of variance in the choice probabilities accounted by the tradeoff model is generally in the nineties. However, the predictions are, on average, off by approximately .05 on a scale from 0 to 1, meaning that there is room for improvement.

The predictions of the tradeoff model are superimposed on the observations in Figures 3 and 4. The tradeoff model reproduces the qualitative predictions that we derived from it in the previous sections: the relative magnitude effect, the hidden-zero effect, the front-end amount effect and its reversal, and the effect of relocating the front-end amount to the back end of both options. However, there are systematic departures from the observations as well. The most prominent anomaly is that the tradeoff model underpredicts the probabilities of choosing *LL* for A_0 = €1,000 (Cells C, F, and I) in the implicit-zero condition. A reparameterization could resolve this, which suggests that the stimulus context had an effect on parameter values, a problem not uncommon in quantitative analyses of intertemporal choice (e.g., Scholten & Read, 2012a).

Finally, to successfully apply the model to the data, we had to upscale "small" differences in the deviation from a uniform distribution between *LL* and *SS* relative to "large" differences (see Appendix). It would thus seem necessary to introduce an appropriate modification to the model in Equation 3. Specifically, the model would come to include *diminishing sensitivity to deviations from a uniform distribution*: The marginal impact of a deviation decreases with its magnitude. We approximated this with a crude binary distinction between small and large differences in deviation between *LL* and *SS*, but, in future applications of the model, some functional form may capture continuously diminishing sensitivity.

[11] Other significant results were an interaction effect between {A_0, A_2} and $A_{0;2}$, $F(4, 1872) = 23.97$, $p < .005$, $\eta_p^2 = .05$, which was a diluted version of the front-end amount effect and its reversal, and two other interaction effects, which were weaker and more subtle: A two-way interaction effect between $A_{0;2}$ and front-end amount versus back-end amount condition, $F(2, 936) = 7.36$, $p < .005$, $\eta_p^2 = .02$, and a three-way interaction effect among {A_0, A_2}, $A_{0;2}$, and front-end amount versus back-end amount condition, $F(4, 1872) = 3.88$, $p < .005$, $\eta_p^2 = .01$.

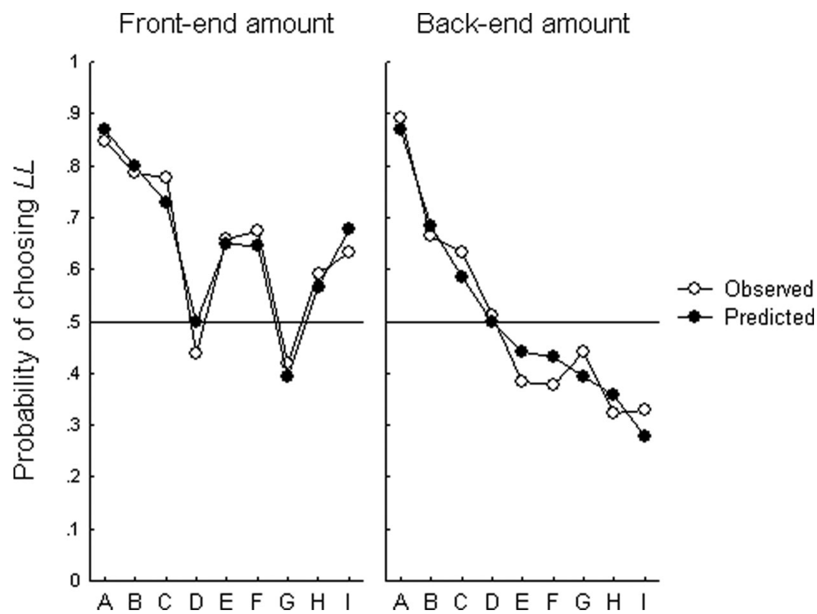

Figure 4. Experiment 3: Observed and predicted probability of choosing *LL* (larger but later) in the nine cells of the within-participant design (from A to I), separately for the front-end amount and back-end amount conditions. The width of the confidence intervals ranged from .061 to .091.

Experiment 4: "Mere" Tokens?

Urminsky and Kivetz (2011) compared the choice between "\$300 in 1 week" and "\$900 in 1 year" with several choices between "A_1 in 1 day and \$300 in 1 week" and "A_1 in 1 day and \$900 in 1 year." Their participants were very impatient, because a large majority declined *LL* in the first choice, where A_1 is an unstated zero outcome. However, when A_1 increased from \$0 to \$10, there was an abrupt increase in the preference for *LL*. As A_1 further increased to \$50, \$100, and \$200, the preference for *LL* increased very little. The authors call this the *"mere" token effect*, because preference is affected by the token but seems insensitive to the magnitude of the token. The mere token effect is a violation of independence in which the common consequence precedes to other consequences.

In the tradeoff model, the mere token effect is the net result of three processes. First, when the token is smaller than the midpoint between the differentiating outcomes (\$300 and \$900), as was the case in Urminsky and Kivetz's (2011) analysis, it introduces a situation in which *LL* deviates more from a uniform distribution than *SS*, which decreases preference for *LL*. Second, the decision maker is less sensitive to a difference between A_1 + \$900 and A_1 + \$300 than to a difference between \$900 and \$300, which also decreases the preference for *LL*. However, the token leads to a much greater decrease in the delay of *LL* (for $q = .3$, by about 3 months) than in the delay of *SS* (no more than about 3 days), and this increases the preference for *LL*.[12] The mere token effect occurs when delay adjustment outweighs preference for spreading and diminishing sensitivity.

The tradeoff model explains why preference was found to be insensitive to the magnitude of the token. First, *LL* deviated more from a uniform distribution than *SS*, but the *difference* in the deviation from a uniform distribution between *LL* and *SS*, that is, $\frac{1}{2}|\$900 - A_1| - \frac{1}{2}|\$300 - A_1$, is independent of A_1 over the range from \$0 to \$300, which includes the narrower range from \$0 to \$200 considered by Urminsky and Kivetz (2011). Second, as shown in Figure 5, sensitivity to the difference between A_1 + \$900 and A_1 + \$300 diminishes very little with A_1 over the range from \$0 to \$200. Third, as shown in Figure 6, the delays of *SS* and *LL* decreased sharply when A_1 increased from \$0 to \$10, but much more slightly as A_1 further increased to \$50, \$100, and \$200. Moreover, had the range of A_1 been extended well beyond \$200, the adjusted delays would have decreased much more sharply again. Most of the action seems to occur between \$50 and \$5,000, which is the range that we explored in Experiment 4.

Method

A total of 349 Portuguese residents (43% male, average age 36 years, 77% having at least completed college or university, and 88% being employed or a student) participated by completing an online questionnaire related to several studies. The present study included three items. Each participant responded to all three. The first item was the tokenless choice between "€200 in 1 week" and "€400 in 1 year." The remaining items, the order of which was randomized across participants, introduced tokens of €50 and €5,000.

[12] Diminishing sensitivity to adjusted delays means that the person is *relatively* less sensitive to the decrease in the longer adjusted delay of *LL* than to the decrease in the shorter adjusted delay of *SS*.

Table 3
Parameter Estimates and Goodness of Fit of the Tradeoff Model

Variable	Description	Experiment 2	Experiment 3
Parameter			
ε	Noise	2.3954	2.6831
γ	Diminishing sensitivity to accumulated outcomes	0.0026	0.0037
σ_U	Preference for spreading of constituent outcomes[a]	0.0023	0.0023
σ_u	Preference for spreading of constituent outcomes[a]	0.0138	0.0441
κ	Tradeoff between time and outcome advantages	6.3775	6.1113
τ	Diminishing sensitivity to (adjusted) delays[b]	0.0000	0.0000
q	Departure from a weighted averaging of delays	0.2995	0.3468
Statistic[c]			
R^2	Goodness of fit	.91	.95
R^2_{adj}	Adjusted goodness-of-fit	.88	.93
RMSD	Badness of fit	.06	.04

[a] Small differences in the deviation from a uniform distribution between larger but longer (*LL*) and smaller but sooner (*SS*) are upscaled relative to large differences (i.e., $\sigma_u > \sigma_U$; see Appendix). [b] Given that the delays were held constant at 0 and 2 and that adjusted delays varied within this narrow range, τ converged to its neutral value of zero (see Appendix). [c] $R^2 = 1 - [\Sigma\,(y - \hat{y})^2/\Sigma\,(y - \bar{y})^2]$, where y is the dependent variable (probability of choosing *LL*), $\hat{y}$ is the predicted value of y, $\bar{y}$ is the mean value of y, and n is the number of data points. $R^2_{adj} = 1 - [\Sigma\,(y - \hat{y})^2/\Sigma\,(y - \bar{y})^2] \times [(n - 1)/(n - k)]$, where k is the number of free parameters, for which this statistic adjusts. RMSD $= \sqrt{\Sigma\,(y - \hat{y})^2/n}$, the root of the mean squared deviation.

Results

In the absence of a token, the choice probabilities were as follows:

Pair A: *SS* Receive €200 in 1 week [54%]
LL Receive €400 in 1 year [46%]

Thus, a small majority preferred *SS*. In the presence of a small token, the choice probabilities were as follows:

Pair B: *SS* Receive €50 tomorrow and €200 in 1 week [50%]
LL Receive €50 tomorrow and €400 in 1 year [50%]

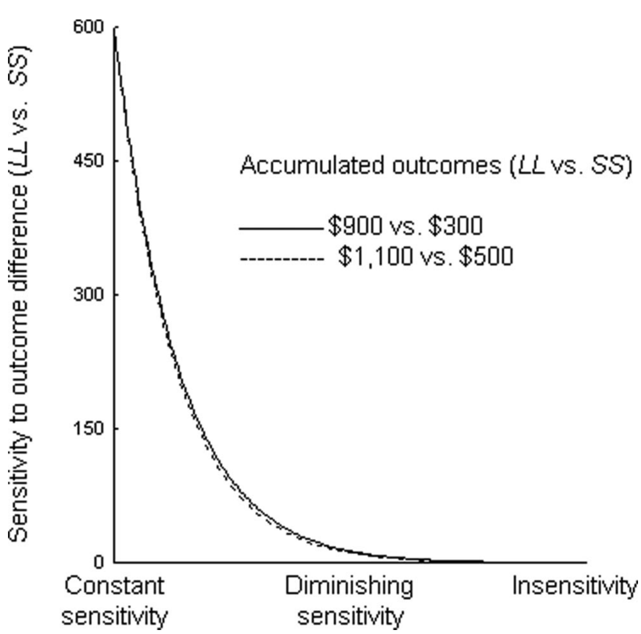

Figure 5. Urminsky and Kivetz's (2011) Experiment 1b: Sensitivity to outcome differences as a function of the accumulated outcomes of *SS* (smaller but sooner) and *LL* (larger but later).

It was a tie between *SS* and *LL*. The increase in the preference for *LL* with the introduction of the small token was marginally significant, $\chi^2(1) = 3.38$, $p < .10$. In the presence of a large token, the choice probabilities were as follows:

Pair C: *SS* Receive €5,000 tomorrow and €200 in 1 week [42%]
LL Receive €5,000 tomorrow and €400 in 1 year [58%]

Thus, a small majority preferred *LL*. The increase in the preference for *LL* with the change from a small to a large token was highly significant, $\chi^2(1) = 10.05$, $p < .005$. We conclude that preference *does* depend on the magnitude of the token, so that the mere token effect is not a "mere" token effect at all.

In the tradeoff model, the greater preference for *LL* in Pair B (the presence of a €50 token) than in Pair A (the absence of a token) is the net result of two processes that hurt *LL* and one process that helps *LL*. First, with the introduction of the token, the sensitivity to the outcome difference between *LL* and *SS* diminishes, which decreases the preference for *LL*, but, as can be seen in Figure 7, it diminishes very little. Second, the token introduces a situation in which *LL* deviates more from a uniform distribution than *SS*, which also decreases the preference for *LL*. However, the token leads to a much greater decrease in the delay of *LL* (for $q = .3$, by about 106 days) than in the delay of *SS* (no more than about 2 days), and this increases the preference for *LL*. The observed pattern shows that the delay adjustment outweighed diminishing sensitivity and preference for spreading.

The greater preference for *LL* in Pair C (the presence of a €5,000 token) than in Pair B (the presence of a €50 token) is the net result of one process that hurts *LL* and two processes that help *LL*. On the one hand, with the greater magnitude of the token, the sensitivity to the outcome difference between *LL* and *SS* diminishes, which decreases the preference for *LL*, and as can be seen in Figure 7, it diminishes visibly. However, the greater magnitude of the token changes a situation in which *LL* deviates more from a uniform distribution than *SS* (175 > 75) into one in which the reverse is true (2400 > 2300), which increases the preference for *LL*. Moreover, the greater magnitude of the token leads to a much greater decrease in the delay of *LL* (for $q = .3$, by about 159 days) than in

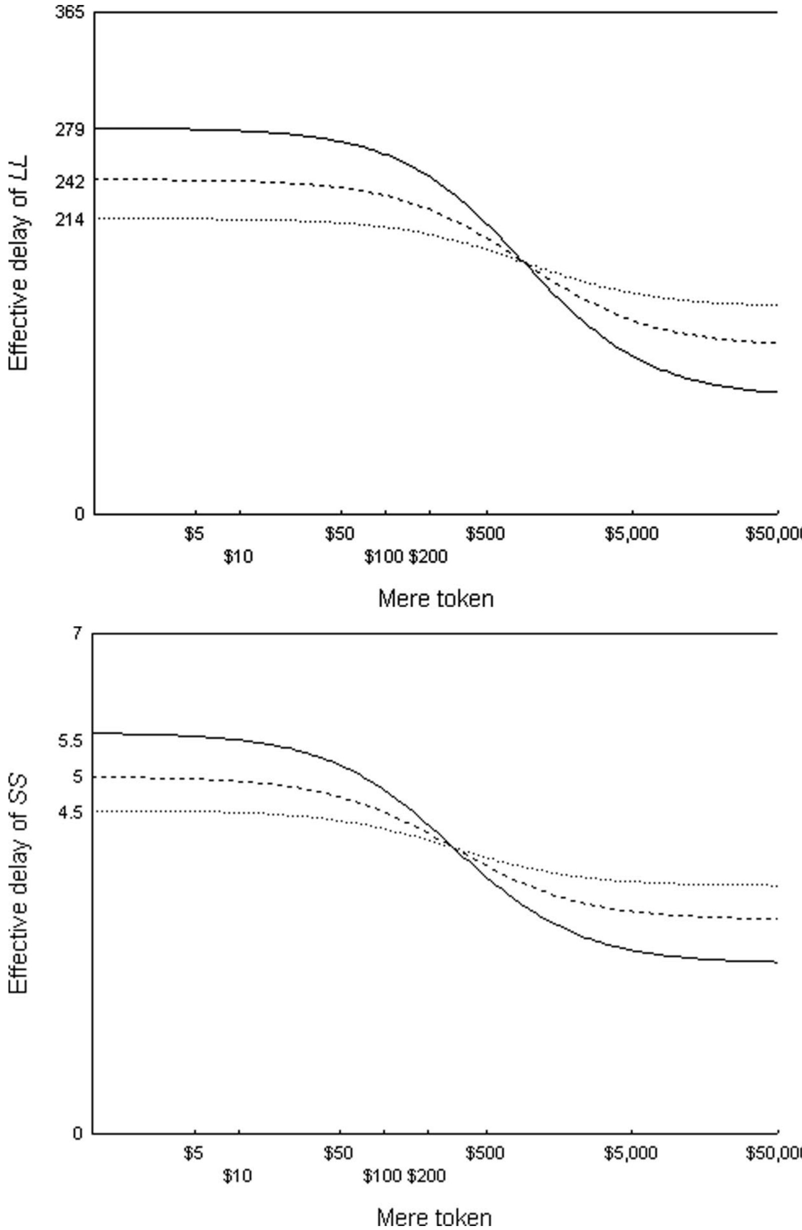

Figure 6. Urminsky and Kivetz's (2011) Experiment 1b: Adjusted delays of *LL* (larger but later; top panel) and *SS* (smaller but sooner; bottom panel) in choices between "A_1 in 1 day and $300 in 1 week" and "$A_1$ in 1 day and $900 in 1 year." Token ($A_1$) is logarithmically scaled. Delays are in days. For the unstated zero token, the delays are 365 (*LL*) and 7 days (*SS*). The departure from weighted averaging is $q = .3$ (solid line), .5 (dashed line), and .7 (dotted line).

the delay of *SS* (no more than about 2.5 days), and this also increases the preference for *LL*. The observed pattern shows that preference for spreading and delay adjustment outweighed diminishing sensitivity.

In sum, the tradeoff model can accommodate both the "mere" token effect and its dependence on the magnitude of the token. Across the four experiments that we conducted, the tradeoff model, as extended to two-outcome sequences, received substantial support. We next discuss some issues raised by our extension of the tradeoff model.

General Discussion

We originally developed the tradeoff model for choices between pairs of single outcomes, a domain in which it accommodates all anomalies that conventional models of intertemporal choice can and cannot address. These models include the discounted utility model (Samuelson, 1937), the (quasi-) hyperbolic discounting model (Laibson, 1997; Loewenstein & Prelec, 1992), and the discounting by intervals model (Scholten & Read, 2006). In this article, we extended the tradeoff model to pairwise choices involving two-outcome sequences. The thrust of our proposal is that

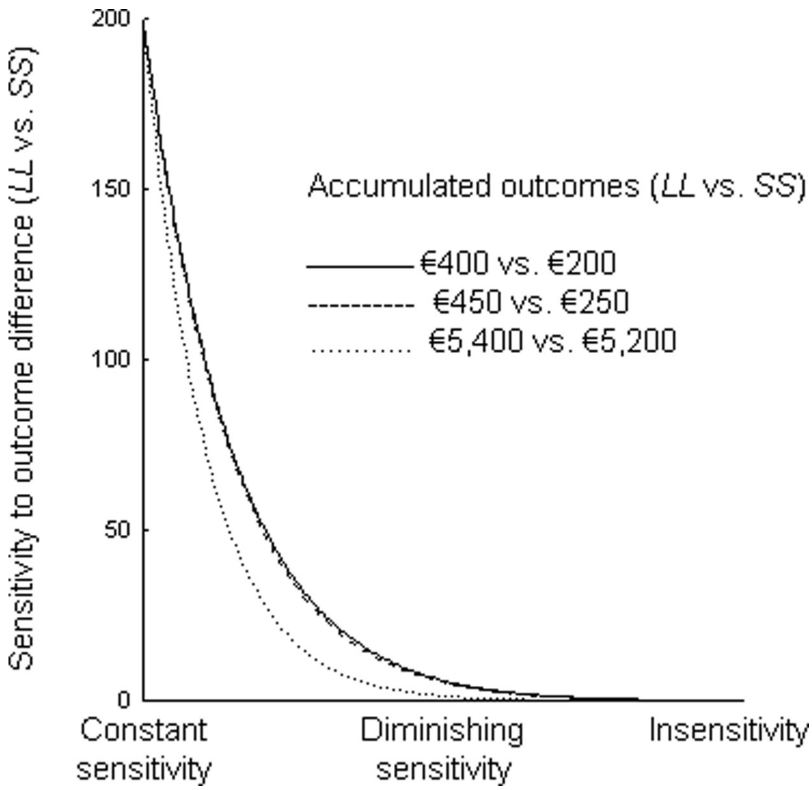

Figure 7. Experiment 4: Sensitivity to outcome differences as a function of the accumulated outcomes of *SS* (smaller but sooner) and *LL* (larger but later).

tradeoffs between sequences are made by weighing accumulated outcomes against outcome-adjusted delays. Thus extended, the tradeoff model accommodates all recently discovered anomalies in pairwise choices involving two-outcome sequences, including the ones documented in this article.

A prominent model of preferences over sequences is the one proposed by Loewenstein and Prelec (1993). In the LP model, the decision maker chooses the sequence with the highest overall value. The overall value of a sequence results from the sum of the (discounted) values of the outcomes in the sequence, preference for improving sequences of outcome values, and preference for uniform sequences of outcome values. Three differences between the LP model and the tradeoff model should be discussed.

First, in the LP model, the outcomes are first valued and then the (discounted) outcome values are summed (value accumulation), whereas, in the tradeoff model, the outcomes are first summed and then the sum is valued (outcome accumulation). In proposing outcome accumulation, we were motivated by what we consider a proper representation of dominance relationships. For instance, the tradeoff model treats the choice between "$500 today" and "$450 today and $50 in 2 months" as a choice between a *dominant* and a *dominated* option, because the accumulated outcomes are the same ($500) while the delay of the dominant option (0 months) is shorter than the adjusted delay of the dominated option (between 0 and 2 months). With value accumulation, however, there would be no dominance relationship: The dominated option would have an advantage over the dominant option along the outcome attribute, because by diminishing sensitivity to outcomes, $v(450) + v(50) > v(500)$. This is prevented by outcome accumulation.

Our step to replace value accumulation by outcome accumulation is remotely similar to the step that Tversky and Kahneman (1992) took in replacing probability weighing by *cumulative* probability weighing. In original prospect theory, *dominance detection* had to be postulated as an editing operation prior to the evaluation of options in order to prevent choice of dominated options. In cumulative prospect theory, this is automatically prevented by cumulative probability weighing.

In proposing outcome accumulation, we were also motivated by the front-end amount effect. In Table 1, preference for *LL* is stronger in Choice A than in Choice B. This, with value accumulation, means that

$$v(30) - v(10) > v(500) + v(30) - v(510),$$

so that

$$v(10) < v(510) - v(500),$$

which is inconsistent with diminishing sensitivity to outcomes (see also Rao & Li, 2011). With outcome accumulation, however, the stronger preference for *LL* in Choice A than in Choice B means that

$$v(30) - v(10) > v(530) - v(510),$$

which is consistent with diminishing sensitivity to accumulated outcomes. Therefore, outcome accumulation, in combination with diminishing sensitivity, yields the front-end amount effect among the countervailing forces of preference for spreading and delay adjustment. In addition, outcome accumulation (first summing outcomes and then valuing the sum) seems cognitively less taxing than value accumulation (first valuing outcomes and then summing the valued outcomes): It requires $k - 1$ fewer operations, where k is the number of outcomes. This may be a motive for evaluating sequences by outcome accumulation rather than value accumulation. If, however, the tradeoff model is to be extended to sequences of qualitative outcomes, outcome accumulation must be revisited.

The second difference between the LP model and the tradeoff model is that the former includes preference for spreading outcome *values*, whereas the latter includes preference for spreading *outcomes*. Preference for spreading outcomes rather than outcome values preserves equal treatment of what contributes to the value of a sequence (the sum of the outcomes) and what detracts from the value of a sequence (the absolute deviation between the outcomes). In addition, evaluating the spread of outcomes requires k fewer operations than evaluating the spread of outcome values. If, however, the tradeoff model is to be extended to sequences of qualitative outcomes, preference for spreading outcomes must also be revisited.

The third difference between the LP model and the tradeoff model is that the former includes preference for improvement, whereas the latter does not. This decision was more pragmatic than principled: We found no evidence of preference for improvement in choices between elementary sequences of monetary outcomes. Moreover, the available evidence on choices between elementary sequences of nonmonetary, qualitative outcomes is equivocal.

In Loewenstein and Prelec's (1993) Example 2, an overwhelming majority preferred visiting an abrasive aunt on one weekend and then visiting friends on the next weekend over visits in reverse order. However, in Loewenstein and Prelec's (1993) Example 1, only a small majority preferred dining at a less liked Greek restaurant in 1 month and then dining at a better liked French restaurant in 2 months over dinners in reverse order, which could be ascribed to a preference for spreading. Moreover, in Frederick and Loewenstein's (2008) Study 2b, an even smaller majority preferred the improving dinner sequence over the declining one, and preference for deterioration emerged when the choice task was replaced by a pricing task. In Frederick and Loewenstein's (2008) Study 2a, preference for deterioration also emerged from a matching task involving monetary two-outcome sequences:

> *A.* Receive $2,000 now for signing up as a participant in a 1-year study and receive another $1,000 when the study is complete.
>
> *B.* Receive $1,000 now for signing up as a participant in a 1-year study and receive another $_____ when the study is complete.

The mean response was well over $2,000, and almost no one gave a response below $2,000. The explanation offered by the tradeoff model is that the adjusted delay of *B* is longer than that of *A*, for which the decision maker must be compensated.

In sum, preference for improvement and, more generally, preferences over sequences depend on the context and method of preference elicitation and in ways that are not yet well understood. The extended tradeoff model may receive its strongest support from choice tasks involving monetary, or, more generally, quantitative, outcomes. In future development of the model, it may be necessary to include preference for improvement.

While we agree that preferences over sequences are conditioned by multiple motives and that different contexts and methods of preference elicitation tap into different motives (Frederick & Loewenstein, 2008), we also consider the current extension of the tradeoff model a promising step in the prediction of preferences over sequences. It covers all the latest anomalies in pairwise choices involving elementary sequences of monetary outcomes. Extension beyond this restricted domain is left as a challenge for future development.

References

Andersen, S., Harrison, G. W., Lau, M. I., & Rutström, E. E. (2010). Behavioral econometrics for psychologists. *Journal of Economic Psychology, 31,* 553–576. doi:10.1016/j.joep.2010.03.017

Frederick, S., & Loewenstein, G. (2008). Conflicting motives in evaluations of sequences. *Journal of Risk and Uncertainty, 37,* 221–235. doi:10.1007/s11166-008-9051-z

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics, 112,* 443–478. doi:10.1162/003355397555253

Leland, J. W. (2002). Similarity judgments and anomalies in intertemporal choice. *Economic Inquiry, 40,* 574–581. doi:10.1093/ei/40.4.574

Loewenstein, G. (1987). Anticipation and the valuation of delayed consumption. *Economic Journal, 97,* 666–684. doi:10.2307/2232929

Loewenstein, G., & Prelec, D. (1992). Anomalies in intertemporal choice: Evidence and an interpretation. *Quarterly Journal of Economics, 107,* 573–597. doi:10.2307/2118482

Loewenstein, G., & Prelec, D. (1993). Preferences for sequences of outcomes. *Psychological Review, 100,* 91–108. doi:10.1037/0033-295X.100.1.91

Lopes, L. L., & Oden, G. C. (1999). The role of aspiration level in risky choice: A comparison of cumulative prospect theory and SP/A theory. *Journal of Mathematical Psychology, 43,* 286–313. doi:10.1006/jmps.1999.1259

Magen, E., Dweck, C. S., & Gross, J. J. (2008). The hidden-zero effect. Representing a single choice as an extended sequence reduces impulsive choice. *Psychological Science, 19,* 648–649. doi:10.1111/j.1467-9280.2008.02137.x

Rao, L. L., & Li, S. (2011). New paradoxes in intertemporal choice. *Judgment and Decision Making, 6,* 122–129.

Read, D. (2001). Is time-discounting hyperbolic or subadditive? *Journal of Risk and Uncertainty, 23,* 5–32. doi:10.1023/A:1011198414683

Read, D., Frederick, S., & Scholten, M. (2012). Outcome framing in intertemporal choice: The DRIFT Model. *Journal of Experimental Psychology: Learning, Memory, and Cognition.* Manuscript submitted for publication.

Roelofsma, P. H. M. P., & Read, D. (2000). Intransitive intertemporal choice. *Journal of Behavioral Decision Making, 13,* 161–177. doi:10.1002/(SICI)1099-0771(200004/06)13:2<161::AID-BDM348>3.0.CO;2-P

Rubinstein, A. (2003). "Economics and psychology"? The case of hyperbolic discounting. *International Economic Review, 44,* 1207–1216. doi:10.1111/1468-2354.t01-1-00106

Samuelson, P. (1937). A note on measurement of utility. *Review of Economic Studies, 4,* 155–161. doi:10.2307/2967612

Scholten, M., & Read, D. (2006). Discounting by intervals: A generalized model of intertemporal choice. *Management Science, 52,* 1424–1436. doi:10.1287/mnsc.1060.0534

Scholten, M., & Read, D. (2010). The psychology of intertemporal tradeoffs. *Psychological Review, 117,* 925–944. doi:10.1037/a0019619

Scholten, M., & Read, D. (2012a). A quantitative comparison between three models of intertemporal choice. *Journal of Risk and Uncertainty.* Manuscript submitted for publication.

Scholten, M., & Read, D. (2012b). Time and outcome framing in intertemporal tradeoffs: The delay-speedup asymmetry. *Journal of Risk and Uncertainty.* Manuscript submitted for publication.

Tversky, A., & Kahneman, D. (1991). Loss aversion in riskless choice: A reference-dependent model. *Quarterly Journal of Economics, 106,* 1039–1061. doi:10.2307/2937956

Urminsky, O., & Kivetz, R. (2011). Scope insensitivity and the "mere token" effect. *Journal of Marketing Research, 48,* 282–295. doi:10.1509/jmkr.48.2.282

Appendix

Quantitative Application of the Tradeoff Model

In this appendix, we provide the details about the estimation of the tradeoff model. The adjusted delays are given by Equation 2, and the tradeoff rule is given by Equation 3.

Outcome Valuation

The value of an accumulated outcome is given by a logarithmic function:

$$v(x_1 + x_2) = \frac{1}{\gamma}\log(1 + \gamma(x_1 + x_2)),$$

where $\gamma > 0$ is diminishing sensitivity. As γ approaches zero, v approaches an identity function, that is, $v(x) = x$ (constant sensitivity); as γ approaches infinity, v approaches a zero function, that is, $v(x) = 0$ for all x (insensitivity). The value function is a logarithmic function, so that the subtractive tradeoff rule in Equations 1 and 3 produces *bilinear interaction effects* (Scholten & Read, 2010): The present value of a delayed outcome increases more steeply with its magnitude over a shorter delay than over a longer one.

Devaluation for the Deviation From a Uniform Distribution: Preference for Spreading

An accumulated outcome is devalued for the deviation of the constituent outcomes from a uniform distribution:

$$v(x_1 + x_2) - \sigma\, d(x_1, x_2),$$

where

$$d(x_1, x_2) = {}^1/_2|x_1 - x_2|.$$

Preference for spreading is captured by $\sigma > 0$. The σ parameter also serves the technical purpose of scaling $d(x_1, x_2)$ relative to $v(x_1 + x_2)$. In the stimulus designs of Experiments 2 and 3, the differences in the deviation from a uniform distribution between *LL* and *SS*, that is, $d(x_{L_1}, x_{L_2}) - d(x_{S_1}, x_{S_2})$, varied substantially across option pairs (in absolute magnitude, from 10 to 1,010), and preliminary analyses showed that it would be necessary to upscale small differences (10, 15, and 20) relative to large differences (between 235 and 1,010). That is, $\sigma_u > \sigma_U$ for small (u) and large (U) differences.

Time Weighing

Like the value of an accumulated outcome, the weight of an adjusted delay is given by a logarithmic function:

$$w(\hat{t}) = \frac{1}{\tau}\log(1 + \tau\hat{t}),$$

where $\tau > 0$ is diminishing sensitivity. As τ approaches zero, w approaches an identity function, that is, $w(\hat{t}) = \hat{t}$ (constant sensitivity). As τ approaches infinity, w approaches a zero function, that is, $w(\hat{t}) = 0$ for all $\hat{t}$ (insensitivity). In the stimulus designs of Experiments 2 and 3, the constituent delays, t_1 and t_2, were held constant at 0 and 2. With $\hat{t}$ varying within this narrow range, τ converged to its neutral value of zero in both experiments.

Choice Rule

The tradeoff model was estimated through a logistic regression, conducted by the solver routine in Microsoft Excel (see also Lopes & Oden, 1999). Specifically, the model was

$$\log(\Omega) = \frac{1}{\varepsilon}[v(x_{L_1} + x_{L_2}) - v(x_{S_1} + x_{S_2}) - \sigma(d(x_{L_1}, x_{L_2}) - d(x_{S_1}, x_{S_2})) - \kappa(w(\hat{t}_L) - w(\hat{t}_S))],$$

where Ω are the odds of choosing *LL*, and $\varepsilon > 0$ is a "noise" parameter (Andersen, Harrison, Lau, & Rutström, 2010). As ε approaches zero, there is no noise, and choice is entirely determined by the model, so that the predicted values of $\log(\Omega)$ go to plus infinity (probability of choosing *LL* is 1) or minus infinity (probability of choosing *LL* is 0). As ε approaches infinity, there is only noise, and choice is not at all determined by the model, so that the predicted values of $\log(\Omega)$ equal 0 (probability of choosing *LL* is ${}^1/_2$). The special case of indifference (Equation 3) arises when the outcome advantage is equal to the time advantage, so that $\log(\Omega) = 0$, regardless of the value of ε.

Received October 5, 2011
Revision received March 14, 2012
Accepted March 19, 2012 ■