

Preferences for Sequences of Outcomes

George F. Loewenstein and Drazen Prelec

Existing models of intertemporal choice normally assume that people are impatient, preferring valuable outcomes sooner rather than later, and that preferences satisfy the formal condition of independence, or separability, which states that the value of a sequence of outcomes equals the sum of the values of its component parts. The authors present empirical results that show both of these assumptions to be false when choices are framed as being between explicitly defined sequences of outcomes. Without a proper sequential context, people may discount isolated outcomes in the conventional manner, but when the sequence context is highlighted, they claim to prefer utility levels that improve over time. The observed violations of additive separability follow, at least in part, from a desire to spread good outcomes evenly over time.

Decisions of importance have delayed consequences. The choice of education, work, spending and saving, exercise, diet, as well as the timing of life events, such as schooling, marriage, and childbearing, all produce costs and benefits that endure over time. Therefore, it is not surprising that the problem of choosing between temporally distributed outcomes has attracted attention in a variety of disciplinary settings, including behavioral psychology, social psychology, decision theory, and economics.

In spite of this disciplinary diversity, empirical research on intertemporal choice has traditionally had a narrow focus. Until a few years ago, virtually all studies of intertemporal choice were concerned with how people evaluate simple prospects consisting of a single outcome obtained at a point in time. The goal was to estimate equations that express the basic relationship between the atemporal value of an outcome and its value when delayed. Although the estimated functional forms would differ from investigation to investigation, there was general agreement on one point: that delayed outcomes are valued less. In economics, this is referred to as “positive time discounting.”

Although plausible at first glance, the uniform imposition of positive discounting on all of one's choices has some disturbing and counterintuitive implications. It implies, for instance, that when faced with a decision about how to schedule a set of outcomes, a person should invariably start with the best outcome, followed by the second best outcome, and so on until the worst outcome is reached at the end. Because nothing restricts the generality of this principle, one should find people preferring a declining rather than an increasing standard of living, deteriorating rather than improving health (again, holding lifetime health constant), and so on.

In the last few years, several studies have independently fo-

cused on this problem and have shown that with choices of this type, people typically exhibit *negative* time preference (i.e., they prefer an improving series of events, with all other things being equal). In this article we present results that confirm the preference for improvement but qualify it in several respects. First, we found that preference for improvement depends on whether a particular choice is viewed by the decision maker as being embedded in a sequence of outcomes. In other words, when the decision frame draws attention to the sequential nature of choice, negative time discounting typically prevails; however, when the frame draws attention to individual components of the choice, positive time preference predominates.

Second, we examined the validity of a common assumption in theoretical treatments of intertemporal choice: that preferences for outcome sequences are based on a simple aggregation of preferences for their individual components. Separable formulations, such as the discounted utility model, predict that the overall value (i.e., utility) of a sequence is equal to the summed values of its component outcomes. The findings we present challenge this prediction. In general, an individual's valuation of complex sequences cannot be extrapolated in a simple way from his or her valuation of components but responds instead to certain “gestalt” properties of the sequence.

Third, we developed and tested empirically a theoretical model of choice over outcome sequences. The model incorporates two motives that are not part of standard discounted utility formulations: a preference for improvement and a desire to spread consumption evenly over time.

In the next section, we present a series of examples of preference patterns that illustrate preference for improvement and preference for spreading good outcomes evenly over time. We then develop a theoretical model of sequential choice that incorporates these two motives. Finally, we present two studies that were designed to test the model parametrically.

Basic Motives Underlying Choices Between Sequences

A temporal sequence is a series of outcomes spaced over time. The outcomes could be specific events, such as one's activi-

George F. Loewenstein, Department of Social and Decision Sciences, Carnegie Mellon University; Drazen Prelec, Department of Psychology, Massachusetts Institute of Technology.

Correspondence concerning this article should be addressed to George F. Loewenstein, Department of Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213.

ties over consecutive weekends, or they could be more abstract economic indexes (e.g., income levels over consecutive years). With a few notable exceptions (Bell, 1977; Epstein & Hynes, 1983; Gilboa, 1989; Horowitz, 1988; Meyer, 1976, 1977), most theoretical treatments of intertemporal choice have been conducted within the framework of the general discounting model, which represents the value of a sequence $X = (x_1, \dots, x_n)$ by the weighted utility formula (Koopmans, 1960; Koopmans, Diamond, & Williamson, 1964; Samuelson, 1937):

$$V(X) = \sum_t w_t u(x_t). \quad (1)$$

The formula implies that whenever two sequences differ in only two periods, then preference between them does not depend on the common outcomes in the remaining $n - 2$ periods (separability). Economic applications normally make two additional assumptions:

1. *Impatience.* The coefficients, w_1 , w_2 , and so forth, are declining, which indicates that earlier periods have greater weight in determining preferences.

2. *Constant discounting.* The marginal rate of utility substitution between any two adjacent periods is the same, $w_{t+1}/w_t = \delta$. This produces the compound discounting formula,

$$V(X) = \sum_t \delta^t u(x_t).$$

In the remainder of this section, we describe simple choice patterns that are inconsistent with these properties.

Preference for Improvement and the Sequence "Frame"

A number of recent studies have shown that people typically favor sequences that improve over time. Loewenstein and Sicherman (1991) found that a majority of subjects preferred an increasing wage profile to a declining or flat one for an otherwise identical job. Varey and Kahneman (in press) studied preferences over short-term streams of discomfort, lasting from 2 to 20 min, and found that subjects strongly preferred streams of decreasing discomfort even when the overall sum of discomfort over the interval was otherwise identical. A preference for experiences that end well has also been documented by Ross and Simonson (1991). In one study, they presented subjects with a series of hypothetical choices between sequences that ended with a loss (e.g., win \$85, then lose \$15) or a gain (lose \$15, then win \$85). Subjects overwhelmingly preferred sequences that ended with a gain.

The preference for improvement appears to depend not only on the amount of improvement but the speed with which it occurs over time—its "velocity"—as Hsee and Abelson (1991) called it (see also Hsee, Abelson, & Salovey, 1991). Subjects in one of their studies played a game in which their probability of winning either decreased or increased over time at one of three rates of change. Those in conditions with increasing probabilities of winning rated the game as more satisfying than those in conditions with decreasing probabilities, and the effect of direction (increase vs. decrease) was amplified by the velocity of the change.

Preference for improvement appears to be an overdetermined phenomenon, driven in part by savoring and dread

(Loewenstein, 1987), adaptation and loss aversion (Kahneman & Tversky, 1979), and recency effects (Miller & Campbell, 1959). Savoring and dread contribute to preference for improvement because, for gains, improving sequences allow decision makers to savor the best outcome until the end of the sequence. With losses, getting undesirable outcomes over with quickly eliminates dread. Although there is evidence that people sometimes like to defer desirable outcomes (Loewenstein, 1987), getting undesirable outcomes over with quickly appears to be more widespread. A number of studies have shown that people prefer immediate rather than delayed electric shocks (Barnes & Barnes, 1964; Carlsmith, 1962). A similar result has been reported by Carson, Horowitz, and Machina (1987) in the context of cigarette smoking. They found that nonsmokers, when asked how much they would need to be paid immediately to smoke a pack of cigarettes either immediately or in 1, 5, or 10 years, specified amounts that increased as a function of time delay.

Adaptation and loss aversion lead to a preference for improving sequences because people tend to adapt to ongoing stimuli over time and to evaluate new stimuli relative to their adaptation level (Helson, 1964). Loss aversion (Kahneman & Tversky, 1979) refers to the observation that people are more sensitive to a loss than to a gain of equal absolute magnitude. It is illustrated by the fact that few people will voluntarily accept a bet that provides an equal chance of winning or losing any given amount. If people adapt to the most recent level of stimuli they experience, then improving sequences will afford a continual series of positive departures (gains) from their adaptation level, whereas declining sequences provide a series of relative losses. Loss aversion implies that the latter will be especially unattractive relative to the former.

The specific psychological mechanisms underlying the adaptation and loss-aversion explanation are somewhat ambiguous. It may be that when faced with a sequence (e.g., a series of increasing or decreasing salary levels), people imagine themselves experiencing the sequence, adapting to the standard of living that each salary level implies, and reacting to negative or positive deviations from such standards. They would then recognize that upward adjustments from one's standard of living are more pleasurable than downward adjustments, leading them to prefer the increasing sequences. Alternatively, adaptation and loss aversion may not involve any explicit anticipation of future experience but may instead be a simple application of perceptual loss aversion. Just as people treat risky outcomes as gains and losses instead of absolute wealth levels, it is possible that they evaluate sequences as series of upward and downward shifts rather than as a series of levels. Loss aversion would then imply that downward shifts receive disproportionate weight. The important difference between these two accounts is that in the former case, the evaluation reflects a type of hedonic forecasting, whereas in the latter case, the preference is instead perceptually driven (i.e., based on the tendency to interpret sequences as gains and losses regardless of how they are actually experienced when they unfold).

The adaptation and loss aversion explanation is closely related to the concept of a "contrast effect" (Elster, 1985; Elster & Loewenstein, in press; Tversky & Griffin, 1991). Contrast effects refer to the effect on one's evaluation of the present of

comparing the present with the past or future. If backward-looking contrast effects are more potent than forward-looking ones, as seems plausible (Prelec & Loewenstein, 1991), then the net impact of contrast effects will be to augment the preference for improvement over time. This is because inferior early experiences will create a favorable contrast that will enhance the utility of later experiences.

A final psychological mechanism that may contribute to the preference for improvement is the recency effect, which has been observed in recall, attitude formation, and belief updating (Miller & Campbell, 1959). As Ross and Simonson (1991) noted, the final outcome in a sequence is likely to be the most salient to the decision maker after the conclusion of the sequence. If decision makers naturally adopt a retrospective perspective when evaluating outcome streams, as Varey and Kahneman (in press) argued, then recency effects will cause late periods to be overweighted relative to those that occur in the middle of the sequence. Likewise, "primacy effects" would promote an overweighting of early periods.

Savoring and dread apply to single-outcome prospects as well as to outcome sequences. For this reason, they can explain why people who otherwise discount the future sometimes defer pleasurable outcomes and get unpleasant outcomes over with quickly. Neither adaptation and loss aversion, nor recency effects, on the other hand, have obvious implications for single-outcome events. Therefore, these latter effects probably do not play a major role in timing preferences for such simple prospects. Adaptation and loss aversion are, however, present as potential factors in outcome sequences. The fact that only one motive for improvement operates for simple outcomes and two operate for sequences suggests that preference for improvement will be stronger in the latter case. Our first example illustrates the preference for improvement and shows that it depends, in part, on whether a particular choice is perceived as being between individual outcomes or sequences.

Ninety-five Harvard University undergraduates were asked the following three questions and were instructed to ignore their own personal scheduling considerations (e.g., preexisting plans) in responding.

Example 1

- | | |
|--|--------|
| 1. Which would you prefer if both were free? | n = 95 |
| A. Dinner at a fancy French restaurant | 86% |
| B. Dinner at a local Greek restaurant | 14% |

For those who prefer French:

- | | |
|--|--------|
| 2. Which would you prefer? | n = 82 |
| C. Dinner at the French restaurant on Friday in 1 month | 80% |
| D. Dinner at the French restaurant on Friday in 2 months | 20% |
| 3. Which would you prefer? | n = 82 |
| E. Dinner at the French restaurant on Friday in 1 month and dinner at the Greek restaurant on Friday in 2 months | 43% |
| F. Dinner at the Greek restaurant on Friday in 1 month and dinner at the French restaurant on Friday in 2 months | 57% |

Because two of the three motives hypothesized to motivate the preference for improvement operate only for sequences of outcomes, we anticipated that a larger fraction of respondents

would prefer to put the fancy French dinner off into the future when it was combined in a sequence with the Greek dinner than when it was expressed as a single-outcome prospect. This was indeed the pattern that we observed. Of the 86% of subjects who preferred the fancy French dinner, 80% preferred a more immediate dinner (Option C) over a more delayed dinner (Option D). Thus, only 20% preferred to delay the French dinner when it was expressed as a single, isolated item. However, when the French dinner was put into a sequence with the Greek dinner, giving subjects the option of having Greek and then French, or French and then Greek, the majority (57%) preferred to defer the French dinner. Even with single-outcome events, there was some motivation to defer the French dinner: Witness the 20% of subjects who opted for the longer delay. However, this tendency was clearly stronger for sequences than for individual items.

We observed the same pattern when we substituted "dinner at home" for the Greek dinner. Because most people eat dinner at home on most nights anyway, the mere embedding of the French dinner in an explicit binary sequence does not introduce any real modification of the problem relative to the single-outcome frame in Question 2. The only thing that happens is that the subject is reminded that the choice is "really" between complete sequences. Like other framing effects, such reminders cause preferences to shift, in this case in favor of the improving sequence.

The pattern of preferences revealed by these choices is incompatible with any discounted utility model, as defined by Equation 1. A preference for a French dinner in 1 month rather than 2 suggests that $w_1 > w_2$; however, a preference for the improving sequence indicates that $w_1u(\text{French}) + w_2u(\text{Greek}) < w_1u(\text{Greek}) + w_2u(\text{French})$, or $w_2 > w_1$, on the assumption that $u(\text{French}) > u(\text{Greek})$, which is confirmed by Question 1.

Defining a Sequence

It appears, then, that two distinct motives are relevant to time preference: impatience and a preference for improvement. Which of these two motives dominates appears to depend on whether the objects of choice are single-outcome prospects or sequences. Impatience dominates choices between single outcomes; the preference for improvement most strongly influences choices between sequences.

In many cases, however, it is not clear whether a particular prospect is properly defined as a sequence. For example, when the attributes of outcomes composing a sequence are incommensurable, or when elements in the sequence are themselves brief but separated by long delays, it seems reasonable to evaluate the elements of the sequence independently of one another. However, when outcomes are commensurable and tightly spaced, the logic for treating them as a sequence will be more compelling. In general, the greater the "integrity" of a series of outcomes, the greater should be its likelihood of being evaluated as an integral sequence.

The following examples illustrate that it is possible to vary the integrity of a sequence so as to influence preferences in a predictable manner. The following three questions were asked of 48 visitors to the Museum of Science and Industry in Chi-

cago. Proportions of subjects giving each response are designated in brackets.

Example 2

Imagine you must schedule two weekend outings to a city where you once lived. You do not plan on visiting the city after these two outings.

You must spend one of these weekends with an irritating, abrasive aunt who is a horrendous cook. The other weekend will be spent visiting former work associates whom you like a lot. From the following pairs, please indicate your preference by checking the appropriate line.

Suppose one outing will take place this coming weekend, the other the weekend after.

	<u>This weekend</u>	<u>Next weekend</u>	
A.	friends	abrasive aunt	[10%] (5/48)
B.	abrasive aunt	friends	[90%] (43/48)

Suppose one outing will take place this coming weekend, the other in 6 months (26 weeks).

	<u>This weekend</u>	<u>26 weeks from now</u>	
A.	friends	abrasive aunt	[48%] (23/48)
B.	abrasive aunt	friends	[52%] (25/48)

Suppose one outing will take place in 6 months (26 weeks from now), the other the weekend after (27 weeks from now).

	<u>26 weeks from now</u>	<u>27 weeks from now</u>	
A.	friends	abrasive aunt	[17%] (8/48)
B.	abrasive aunt	friends	[83%] (40/48)

In the first question, the series of outcomes unfold over a fairly short period (2 weeks) so that we would expect discounting to be relatively weak and the preference for improving sequences to be strong. Here, 90% of subjects opted for the improving sequence. In the second set of options, the absolute interval is much longer (26 weeks), reducing the integrity of the sequence. Here, we would expect discounting to have a greater impact relative to the preference for improvement. Indeed, a much smaller fraction of subjects (52%) chose the improving sequence given the long absolute delay. In the third pair, the sequence interval was once again reduced to 1 week, so we would anticipate a greater preference for the increasing sequences than found in the second pair. However, intuitively, we expected that the long delay prior to the beginning of the sequence would reduce its integrity to some degree. This may explain the slight reduction, relative to the first set of alternatives, in the fraction of subjects opting to get the unpleasant visit over with quickly.

It is not possible to interpret the three modal choice patterns in terms of conventional time preference. The first and third question indicates a negative rate of time preference (i.e., $w_0 < w_1$ and $w_{26} < w_{27}$) in the context of a discounted utility model (see Equation 1). It is safe to generalize that the aunt would be scheduled in the earlier of any consecutive two weekends (i.e., $w_t < w_{t+1}$). Yet, the transitive conclusion does not follow because the middle question implies that $w_0 > w_{26}$. We call this pattern the “magnet effect” because preferences for the two outcomes resemble the behavior of magnets. When distant from one another, two magnets interact only weakly; however, when

brought into close proximity, they exert a force on one another, which causes them to reverse position.

Preference for Spreading

In addition to the desire for improvement over time, preferences also indicate a sensitivity to certain global or “gestalt” properties of sequences having to do with how evenly the good and bad outcomes are arranged over the total time interval. Unlike the question of positive versus negative time preference, which by now has received some attention, there have been few efforts to examine how people like to distribute outcomes over time. Important exceptions are two recent studies examining whether people like to experience two positive, negative, or mixed (positive and negative) events on the same or on different days. Applications of prospect theory imply that people should like to spread gains out across different days and to concentrate losses in a single day, and, although the first of these predictions is generally supported, the evidence for the latter is far more tenuous. Thaler and Johnson (1990) found that people generally expected to be happier when two gains (e.g., winning \$25 in an office lottery and winning \$50 in another) were separated by an interval but also expected to be less unhappy when two losses were separated. Thaler and Johnson argued that prior losses may sensitize people to subsequent losses, contrary to the prediction of prospect theory. Linville and Fischer (1991) likewise failed to observe a concentration of losses, although the tendency to do so was greater for small losses than for large ones. They explained this shift in terms a model of “coping capacity,” which postulates that people have a limited psychological capacity to absorb losses and that they may wish to separate losses in time in order to replenish their coping resources. Thus, for both gains and losses, there does appear to be a preference for spreading outcomes out over time.

Although these results are suggestive, their applicability to the types of sequences we were concerned with is limited. First, all of the choices in these studies involved only two outcomes. This made it possible to examine whether people like to concentrate or spread out outcomes, but not whether they exhibit more complex patterns of preference (e.g., for certain types of patterns of outcomes over time). Second, the choices were all between experiencing outcomes at the same or different points in time. Such a design leaves unanswered, for those who prefer to separate outcomes, the question of how much of a gap is ideal.

The central insight that we took from this research was that, when presented with more than one same-valence outcome, people generally like to spread outcomes over time rather than concentrating them. The following problem presented to 37 Yale University undergraduates (from Loewenstein, 1987) illustrates this desire for spread. Subjects were first given a choice between Options A and B, then between Options C and D; they were instructed to ignore scheduling considerations. Percentages who chose each of the options are presented in the right-hand column.

Example 3

Which would you prefer?

Option	This weekend	Next weekend	Two weekends from now	Choices
A	Fancy French	Eat at home	Eat at home	16%
B	Eat at home	Fancy French	Eat at home	84%
C	Fancy French	Eat at home	Fancy lobster	54%
D	Eat at home	Fancy French	Fancy lobster	46%

Choosing between Options A and B, the majority of subjects preferred to postpone the fancy dinner until the second weekend, consistent with the widespread preference for improvement. However, the insertion of the common lobster dinner in Options C and D caused preference to shift slightly in favor of having the French dinner right away. This pattern violates additive separability (and any model given by Equation 1). Because the third period is identical for Options A and B and for C and D, separability implies that anyone who prefers A (B) should prefer C (D).

We believe that the relative attractiveness of Options B and C stemmed in part from the fact that they “covered” the 3-week interval better than did their alternatives. Option A exposed the decision maker to a 2-week period of eating at home, whereas Option B placed the one pleasurable event at the center of the interval. Option D concentrated all of the pleasure at one extreme of the 3-week period, whereas Option C distributed the fancy dinners more evenly over time.

It is worth noting that loss aversion as traditionally conceived would not predict this type of violation. Of the four sequences in the example, only Option D was strictly increasing. It appears, therefore, that a person who strongly dislikes utility reductions across adjacent periods will have a greater tendency to prefer Option D over C than Option B over A, which is the opposite of the observed pattern of choice.

Another possibility is that people have a net liking for changes in utility between adjacent periods, as permitted in Gilboa’s (1989) model. In that case, Option B has more between-period variation than Option A, and Option C more variation than Option D, in accord with the modal preferences. To rule out this explanation, and to show that preferential interactions occur between nonadjacent periods, we modified the original example by inserting additional “eat at home” weekends between the original first two weekends and between the second and third weekends (see Example 4).

Example 4

Imagine that over the next five weekends you must decide how to spend your Saturday nights. From each pair of sequences of dinners below circle the one you would prefer. “Fancy French” refers to dinner at a fancy French restaurant. “Fancy lobster” refers to an exquisite lobster dinner at a four-star restaurant. Ignore scheduling considerations (e.g., your current plans).

Option	First weekend	Second weekend	Third weekend	Fourth weekend	Fifth weekend
A	Fancy French	Eat at home	Eat at home	Eat at home	Eat at home
B	Eat at home	Eat at home	Fancy French	Eat at home	Eat at home
C	Fancy French	Eat at home	Eat at home	Eat at home	Fancy lobster
D	Eat at home	Eat at home	Fancy French	Eat at home	Fancy lobster

In this example, it is not possible to produce an independence violation solely on the basis of adjacent periods. Suppose, for example, that the utility function for sequences has a special set of functions, f_t , that register the impact of adjacent utility levels on preferences:

$$V[(x_1, \dots, x_n)] = \sum_{t=1}^n u_t(x_t) + \sum_{t=1}^{n-1} f_t(x_t, x_{t+1}).$$

In that case, the difference in value between Sequences A and B still equals the difference in value between Sequences C and D because

$$\begin{aligned} V(A) - V(B) &= u_1(F) + u_3(H) + f_1(F, H) + f_2(H, H) \\ &\quad + f_3(H, H) - u_1(H) - u_3(F) - f_1(H, H) \\ &\quad - f_2(H, F) - f_3(F, H) \\ &= V(C) - V(D), \end{aligned}$$

where F, H refer to the French and home dinners. In the survey, visitors to the Museum of Science and Industry in Chicago ($N = 51$) strongly preferred Sequence B to A (88%) and just slightly preferred Sequence C to D (51%). Despite the different subject population (museum visitors vs. undergraduates), and despite the inclusion of two “filler” weekends, the preference pattern was virtually unchanged from Example 3. The violation of independence cannot therefore be attributed to particular feelings about utility changes from one period to the next.

Marking an Interval

How one distributes events over an interval clearly depends on the duration of that interval. However, in the real world the relevant interval is often ambiguous and may vary for different people and for different types of outcomes. For example, the relevant interval for a free dinner is probably shorter than for a free round-trip flight. The final example we present was constructed to test whether manipulation of the implicit interval would influence timing preferences. One hundred one visitors to Chicago’s Museum of Science and Industry were asked to schedule two hypothetical free dinners at the restaurant of their choice.

Example 5

Suppose you were given two coupons for fancy dinners for two at the restaurant of your choice. The coupons are worth up to \$100 each. When would you choose to use them? Please ignore considerations such as holidays, birthdays, etc.

One third of the subjects were asked to schedule the two dinners without any imposed time constraints (unconstrained group). Another one third, the “4-month constraint group,” were told that they could use the coupons any time in the next 4 months. The remaining third, the 2-year constraint group, were told “You can use the coupons at any time between today and 2 years from today.” Although the constraints seemingly *limited* the subjects’ abilities to delay the dinners, we felt that constrained subjects would prefer to delay the dinners more than would unconstrained subjects because the default interval for unconstrained subjects was actually shorter than the explicit intervals faced by the constrained subjects. Table 1 presents our findings.

Table 1
Chosen Delay for First and Second Dinner

Condition	First dinner	Second dinner
Unconstrained		
<i>M</i>	3.3	13.1
<i>Mdn</i>	2.0	8.0
Four-month constraint		
<i>M</i>	3.0	10.4
<i>Mdn</i>	2.0	12.0
Two-year constraint		
<i>M</i>	7.7	31.1
<i>Mdn</i>	4.0	26.0

Medians were a more representative measure of population preferences here because they attenuated the impact of a few extremely long delays obtained in the unconstrained condition. Considering the medians, the median delay interval for the second dinner was longer under either constrained condition than the median delay in the unconstrained condition. The effect was especially prominent, however, for the 2-year delay, suggesting that the average default interval for subjects might have been close to 4 months. Overall, the effect of the two constraints on the selected delays was significant, according to a one-way analysis of variance (ANOVA), $F(2, 96) = 3.7$, $p < .03$, for the first dinner, and $F(2, 96) = 13.2$, $p < .0001$, for the second.

These results are inconsistent with the economic axiom of revealed preference, according to which the imposition of a time constraint on an initially unconstrained population should affect only the responses of that fraction of the population whose preferred delays are longer than permitted by the constraint. Therefore, the population averages should be longer in the unconstrained condition.

A Model for Preferences Over Outcome Sequences

Taken together, these examples confirm that subjects may or may not frame a given intertemporal choice as one involving sequences of outcomes and that whether they do so can significantly influence their behavior. Furthermore, when people do see themselves as choosing between sequences, their preferences are poorly captured by conventional discounting models.

In this section, we develop a nonseparable model that accommodates the anomalous preference patterns presented in the previous section. The guiding idea behind the model is that evaluation of sequences reflects the interaction between two motives: a basic preference for improvement tempered by a desire to spread the better outcomes more or less uniformly over the entire interval.

In the most general sense, these assumptions are not new. A number of recent nonseparable models have added the rate of utility *change*—either discrete (Gilboa, 1989) or continuous (Bordley, 1986; Frank, 1989, in press; Hsee & Abelson, 1991)—as a contributing factor for preference. The distinctive aspect of our approach is that the notions of utility improvement and uniformness are defined with respect to global rather than local sequence properties.

The term *global*, as we use it, refers to comparisons at a given point in time between the set of outcomes that are yet to occur and the set of outcomes that have already occurred. A globally improving sequence would be one in which, at any point in time, the average utility of the remaining periods is greater than the average utility of the periods that have gone by. For example, the sequence (4,2,3,1,2,0) undoubtedly creates a feeling of global decline, even though it includes two increasing transitions along the way. We feel that the enjoyment of the good initial periods, as well as of the positive transitions between Periods 2 and 3, or 4 and 5, would be inhibited by the deteriorating overall pattern.

Similarly, a globally uniform sequence is one in which, at any point in time, the future offers approximately equal average utility as did the past, notwithstanding local fluctuations in utility levels. We posit that such even spreading of good outcomes contributes to the attractiveness of a sequence, with other things being equal.

By contrast, nonseparable rate-of-change models would identify improvement with a positive change in utility between one period and the next and smoothness with small utility fluctuations between adjacent periods. The distinction between the local and global smoothness (or uniformness), for instance, is illustrated by the sequences (0,0,1,1,1,0,0) and (0,1,0,1,0,1,0). The first sequence is more smooth in the local sense because there are only two points (rather than six) at which adjacent utility levels differ; the second sequence achieves a more uniform pattern of outcomes in the global sense, in that the differences between average future utility and average past utility are smaller.

The Decumulated Utility Graph and the Reference Line

The relevant global properties to which we have been referring are displayed visually in a *decumulated* utility graph (see Figure 1), which plots total remaining utility as function of period number. This picture can be interpreted as a “utility-budgeting” chart, showing how much utility has been spent up to a point and how much is left over for subsequent periods. The particular pattern in Figure 1 corresponds to a utility series that starts low, peaks around the middle of the interval, and then declines toward the end.

In a decumulated utility graph, a constant series of utility levels is represented by a negatively sloped straight line. Taking our cue from the well-documented pervasiveness of reference point effects in many different choice domains, we postulate that the decumulated graph of the uniform sequence is the natural global reference point concept for the sequence domain. More precisely, the internal *reference line* for each sequence is the graph of the unique constant sequence that gives the same total utility as the original one.

When the decumulated graph falls below the reference line, utility is being used up at a faster-than-average pace, and leaner periods will necessarily follow; above the reference line, utility is being saved up, and future periods will improve on the previous ones.

In the next section, we develop the formal definition of our model through an accounting scheme for utility derived from anticipation. This simple model captures the essence of the




Figure 1. The decumulated utility graph showing remaining utility as a function of time. (The corresponding utility sequence would be bell-shaped.)

improvement motive, and the motive for uniform spreading of outcomes, and can therefore explain most of the qualitative preference patterns presented earlier. We then show how the separate impact of time discounting can be incorporated into a more generalized version of the model.

Definition of Improvement and Uniformness

It is commonplace that events create satisfaction or distress not just while they take place but also before they occur, through anticipation, and after they occur, through recollection. The simplest possible accounting of these two additional sources of value would have the total *anticipated* utility computed as a sum in which each utility level u_t is multiplied by the number of periods that precede it ($t - 1$),

$$AU = \sum_t (t - 1)u_t,$$

and total *recollected* utility as a sum in which each utility level is multiplied by the number of subsequent periods ($n - t$),

$$RU = \sum_t (n - t)u_t.$$

For instance, if a positive event is scheduled on Day 4 of a 5-day sequence, this will yield 3 days of anticipation and 1 day of recollection. (The measures of anticipation and recollection, of course, depend on the framing of the problem, in this case the number of days under consideration.)

In the decumulated utility graph (see Figure 1), anticipated utility is represented by the area to the left and below the curved line, and recollected utility is represented by the area to the right and above the line. The value of the cumulative utility

distribution at any point indicates how much of the sequence's overall utility remains to be experienced. We now define the net improvement that a sequence provides as $(AU - RU)/2$, or one half of the excess of anticipated utility over recollected utility. The $1/2$ factor makes it possible to read the improvement measure directly from the decumulated graph: It is the difference between the area where the decumulated line falls above the reference line (labeled "positive deviation" in Figure 2) and the area where it falls below it (labeled "negative deviation" in Figure 2).

With regard to the second of the two basic motives, we define the deviation from uniform utility spreading in a particular sequence as the absolute difference between the decumulated graph for that sequence and its reference line, which is to say, the sum of the two shaded areas in Figure 2.

The Model

Expressing the relevant areas formally involves some algebra. For each period t , we first compute the difference, d_t , between the cumulated utility received up to that period and the cumulated utility that *should have been received* had the utility total been allocated in a perfectly uniform manner across the n periods:

$$d_t = \frac{t}{n} \sum_{i=1}^n u_i - \sum_{i=1}^t u_i. \tag{2}$$

Geometrically, the value of d_t equals the vertical difference between the decumulated utility line and the reference line. To confirm that the aggregate of all d_t s indeed equals net improvement, we calculate that

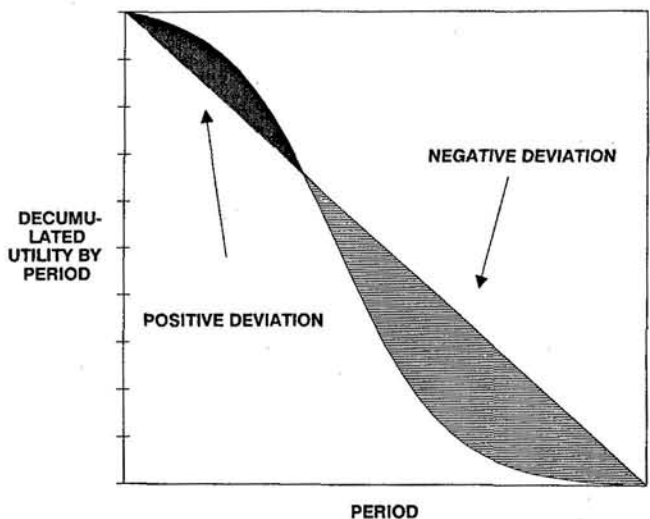


Figure 2. Positive and negative deviations defined as the distance between the decumulated utility function and the negatively sloping reference line.

$$\begin{aligned}
\sum_{t=1}^n d_t &= \sum_{j=1}^n \left(\frac{t}{n} \sum_{i=1}^n u_i - \sum_{i=1}^t u_i \right) \\
&= \frac{\sum_{t=1}^n t}{n} \sum_{i=1}^n u_i - \sum_{t=1}^n \sum_{i=1}^t u_i \\
&= \frac{n+1}{2} \sum_{i=1}^n u_i - \sum_{t=1}^n (n+1-t) u_t \\
&= \frac{1}{2} \left[\sum_{t=1}^n (t-1) u_t - \sum_{t=1}^n (n-t) u_t \right] \\
&= \frac{1}{2} (AU - RU).
\end{aligned}$$

The terms d_t allocate the difference $(AU - RU)/2$ to individual periods, so that positive values are associated with periods at which the future is better than the past and negative values to periods when it is worse than the past.

The value of a sequence is now defined as a weighted sum of three factors: the total utility provided by the sequence, the net improvement, and the deviation from uniform utility spreading:

$$\text{Value} = \sum_{t=1}^n u_t + \beta \sum_{t=1}^n d_t + \sigma \sum_{t=1}^n |d_t|. \quad (3)$$

Interpretation of the Parameters

The parameter β signals whether a person prefers improving ($\beta > 0$) or declining ($\beta < 0$) sequences, whereas σ determines whether he or she prefers uniform ($\sigma < 0$) or nonuniform sequences ($\sigma > 0$). This is a straightforward interpretation of the parameters. However, an alternative and psychologically suggestive set of interpretations can be obtained through a more explicit articulation of the reference point aspect of the model.

We first split the deviation terms into their positive and negative parts:

$$d_t^+ \equiv \frac{|d_t| + d_t}{2}, \quad d_t^- \equiv \frac{|d_t| - d_t}{2}. \quad (4)$$

In words, d_t^+ equals d_t if it is positive, and zero otherwise; d_t^- equals $|d_t|$ if it is negative, and zero otherwise (note that $d_t = d_t^+ - d_t^-$). This allows us to rewrite the model so that the net impact of improvement and deterioration on sequence value is assessed separately:

$$\text{Value} = \sum_{t=1}^n u_t + (\beta + \sigma) \sum_{t=1}^n d_t^+ + (\beta - \sigma) \sum_{t=1}^n d_t^-. \quad (5)$$

If indeed the reference line functions as a reference point, and if improvement (d_t^+) defines the attractive direction away from the reference point, then we would predict the following: $\beta > -\sigma > 0$, that is, $(\beta - \sigma) > (\beta + \sigma) > 0$. The desire for uniform utility levels can be viewed as a consequence of loss aversion, which, in the context of our global variables, implies that the dislike of global deterioration, measured by the coefficient $(\beta - \sigma)$, is stronger than the liking for global improvement, measured by $(\beta + \sigma)$. Without any further assumptions, such global

loss aversion would create a motive for interleaving good and bad outcomes more or less evenly over time.

This interpretation covers only one of the eight ordinaly distinct sign-magnitude combinations for β and σ . The remaining possibilities have other interpretations, suggested by the labels in Figure 3. The center of the figure is the point $\beta = \sigma = 0$, identifying a person who compares sequences only by total utility; beyond that, all patterns are equally good. The horizontal axis in the figure represents the β coefficient; points along that line identify individuals who have intertemporally additive preferences, favoring either the present ($\beta < 0$) or the future ($\beta > 0$). Points along the vertical axis, which plots the value of the σ coefficient, identify people without a time orientation who are concerned only with even distribution of utility.

Three of the four corners in Figure 3 have intuitive interpretations. The lower left corner corresponds to the pattern $\beta = \sigma < 0$, which indicates a dislike of improvement but indifference to decline. This describes *impatience* in the nontechnical sense of the term, which is to say, a specific aversion to waiting for good outcomes. The bottom right corner is the pattern $\beta = -\sigma > 0$, which describes the preferences of a person who is indifferent to improvement (as the coefficient for $\sum d_t^+$, $\beta + \sigma = 0$) but is averse to decline; such a person would sacrifice some total utility in order to avoid a declining sequence but would not sacrifice any utility to transform a constant sequence into an increasing one. We label this *dread orientation*. Finally, the complementary pattern is found in the top right corner, where $\beta = \sigma > 0$. This identifies a person who likes improvement but is indifferent to decline (again, in the sense of being willing to sacrifice total utility for one or the other pattern). Such a person derives additional satisfaction from future good outcomes but is not adversely affected by future negative outcomes. A *savoring orientation* would seem to describe this fortunate combination.

Predictions of the Model and a Generalization

Separability Violations

Example 3 induced a violation of independence having the following schematic form: $(1,0,0) < (0,1,0)$, but $(1,0,2) > (0,1,2)$, where, 0 = "eat at home," 1 = "fancy French," and 2 = "fancy lobster," as in Loewenstein (1987). We first calculate the deviations (d_1, d_2) for the four sequences:

$$(1,0,0): d_1 = .33 - 1.0 = -.67; \quad d_2 = .67 - 1.0 = -.33;$$

$$(0,1,0): d_1 = .33 - 0.0 = +.33; \quad d_2 = .67 - 1.0 = -.33;$$

$$(1,0,2): d_1 = 1.0 - 1.0 = 0.0; \quad d_2 = 2.0 - 1.0 = +1.0;$$

$$(0,1,2): d_1 = 1.0 - 0.0 = +1.0; \quad d_2 = 2.0 - 1.0 = +1.0.$$

The value of each sequence is then obtained by applying Equation 3:

$$V[(1,0,0)] = 1 - \beta + \sigma;$$

$$V[(0,1,0)] = 1 + \frac{2}{3}\sigma;$$

$$V[(1,0,2)] = 3 + \beta + \sigma;$$

$$V[(0,1,2)] = 3 + 2\beta + 2\sigma.$$

Because σ is negative, it is possible for a person to prefer sequence $(0,1,0)$ to $(1,0,0)$ and sequence $(1,0,2)$ to $(0,1,2)$:

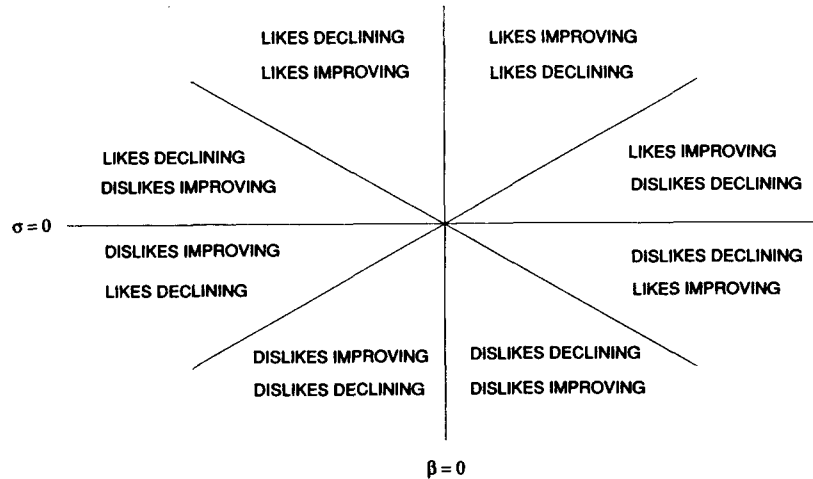


Figure 3. Partitioning of the (β, σ) parameter space into the eight possible sign-magnitude combinations. (The pair of labels in each segment identifies the major [top] and minor [bottom] motive associated with parameter values in that segment.)

$$V[(1,0,0)] - V[(0,1,0)] = -\beta + .33\sigma < 0,$$

$$V[(1,0,2)] - V[(0,1,2)] = -\beta - \sigma > 0.$$

This formal explanation can be interpreted in two ways, which cannot be discriminated at the level of our model. First, it is possible that a person is fundamentally concerned about even spreading, in which case the sequence $(0,1,0)$ is more uniformly spread than $(1,0,0)$, but $(0,1,2)$ is not as uniformly spread as $(1,0,2)$. The other interpretation would attribute the reversal to a fundamental aversion to declining average utility levels (i.e., negative d_i values). In the first pair, the sequence $(1,0,0)$ is globally declining and is possibly rejected on those grounds alone. In the second pair, however, *both* sequences are globally increasing, so that a person is free to indulge the initial period with the $(1,0,2)$ sequence without disturbing the sense of global improvement.

Scheduling Decisions

The model makes definite predictions about the optimal scheduling of an enjoyable event, when there are n consecutive and otherwise indistinguishable dates available. In general, the optimal period will depend on all of the parameters in the model. However, the net effect of the spreading motive will be to move this period closer to the middle of the interval.

Because utility is represented by an interval scale (see the next section), we can assume without loss of generality that the single enjoyable event in an n -period sequence yields a utility level of exactly n and the background events utility zero. The uniform reference sequence then has a cumulated utility distribution $(1, 2, 3, \dots, n)$. By scheduling the event in period t^* , one creates deviations $d_i = +t$ for the periods prior to t^* and $d_i = n - t$ for periods later than (and including) t^* . The sum of absolute deviations

$$[\sum_{t < t^*} t + \sum_{t \geq t^*} (n - t)]$$

is minimized at $t^* = n/2$, the midpoint of the interval.

We feel that this explains the otherwise puzzling results revealed by Example 5. Recall that subjects who were unconstrained in the scheduling of a free meal chose to schedule it sooner, on average, than subjects who were constrained by a long interval. If the explicit interval (2 years) is longer than the implicit one used by unconstrained subjects, then the attraction toward the respective midpoints should further stretch out the optimal scheduling times for the explicitly constrained subjects.

Discounting and the General Model

The model we have just defined represents a radical simplification of intertemporal preferences: Time enters into the picture only in establishing the ordering of events. Whether this is a good approximation to preferences will probably depend on the duration of the entire sequence (the planning interval), as well as on its temporal proximity. Other things being equal, we expect that shorter planning intervals will improve the approximation, whereas temporal proximity will reduce it through a disproportionate weighting of the initial periods.

In this section, we show that the basic model can be readily generalized to incorporate time discounting and other types of differential weighting patterns. We take as a benchmark the most general discounted utility model, restated as follows:

$$V(X) = \sum_t w_t u_t.$$

The model we now introduce retains an additive structure, but, in keeping with the reference point framework developed earlier, it assumes that the sequence value is determined by the

positive and negative deviations of the decumulated utility graph from the reference line:

$$V(X) = \sum_{t=1}^n u_t + \sum_{t=1}^{n-1} \sigma_t^+ d_t^+ + \sum_{t=1}^{n-1} \sigma_t^- d_t^-. \quad (6)$$

In words, the value of any sequence is a separately weighted sum of the positive and negative deviations from the reference line. As in the basic model, the first term gives weight to the total value of outcomes in the sequence; it cancels out if the sequences being compared are permutations of the same set of outcomes. The remaining two terms add up the psychic costs and benefits of facing rising (d_t^+) or falling (d_t^-) utility levels; what is new here is that these costs and benefits are allowed to differ across periods instead of being applied to the summations Σd_t^+ and Σd_t^- .

This model can also be written in a way that brings out more clearly the relation to the discounting formula:

$$V(X) = \sum_{t=1}^n w_t u_t + \sum_{t=1}^{n-1} \sigma_t |d_t|. \quad (7)$$

The first term is the general discounting formula (Equation 1), and the second is the weighted sum of absolute deviations of the cumulated graph from the reference line. The weights, σ_t , are averages of the positive and negative coefficients in Equation 6:

$$\sigma_t = \frac{\sigma_t^+ + \sigma_t^-}{2}, \quad (8)$$

whereas the discounting weights, w_t , have a more complicated relation to the coefficients in Equation 6 (see the Appendix). In the basic model (Equation 3), the coefficients, σ_t , were all identical and presumed negative; the weights, w_t , were linear in period number and presumed increasing, in line with the preference for improvement:

$$w_t = \left(1 + \frac{\beta(n+1)}{2}\right) + \beta t. \quad (9)$$

Discounting enters into the picture when the periods become sufficiently separated in real time or when the first period is close to the present. There is great flexibility in how one could model this through the w_t coefficients. We propose that the relation between (w_1, \dots, w_n) and the real-time parameters of the sequence (τ_1, \dots, τ_n) be modeled in the following way:

$$\begin{aligned} w_t &= f(\tau_t)(\alpha + \beta t), \\ \sigma_t &= f(\tau_t)\sigma, \end{aligned} \quad (10)$$

where $f(\tau_t)$ is the discount function derived empirically from single-outcome intertemporal choices and applied to the time of the t -th period, τ_t , whereas α , β , and σ play the same role as in the earlier model.

The interplay between discounting and improvement can be seen in Figure 4, which shows how the weights for a five-period sequence depend on the real-time parameters of the sequence. The solid line in the figure is the underlying discount function, $f(\tau)$. The weights for the distant, closely spaced sequence are increasing and nearly linear because the discount function values $f(\tau_1)$ through $f(\tau_5)$ are virtually constant over the five dates; the desire for improvement, which depends on period




Figure 4. Individual period weights (w_t) for three 5-period sequences that arise when a linear preference for improvement is modulated by a hyperbolic discount function (solid line).

position rather than absolute time, dominates the discounting effect. Moving that same sequence closer in time to a point at which the hyperbolic discount function is steep produces significant within-sequence discounting, which combines with the linear preference for improvement to produce a U-shaped weighting pattern. Finally, for the widely spaced sequence discounting predominates, and the weights essentially replicate the declining shape of the discount function over time.

This pattern of weights would be consistent with the preference reversal documented in Example 2, in which a visit to an unpleasant relative had to be scheduled in one of two weekends. For adjacent weekends, the values of the discount function at the two weekend dates were not different enough to overwhelm the preference for improvement; however, if the weekends were separated by 6 months, then real-time discounting made it attractive to schedule visiting the aunt in the later slot.

Three Properties of the General Model

Notwithstanding the additional complexity, the general model has a number of properties that are either identical or related to the discounting formula.

1. *Interval utility scale.* As in the general discounting formula, preferences are invariant under linear, positive transformations of the utility function (see the Appendix).

2. *Linearity in the small.* Small improvements in the utility of an outcome have a linear impact on sequence value, as if the model were separable. This considerably simplifies the evaluation of certain common scheduling decisions. Consider, for example, whether it is better to schedule some event in period t or in period $t+1$. In a separable model, this would depend only on the discounting weights w_t and w_{t+1} insofar as it would be desir-

able to redistribute utility from period t to period $t + 1$ whenever $w_t < w_{t+1}$:

$$\frac{\partial V}{\partial u_t} - \frac{\partial V}{\partial u_{t+1}} = w_t - w_{t+1} < 0. \quad (11)$$

In our model, one also has to take into account the impact of any such redistribution on the shape of the cumulated utility graph. However, when the redistribution involves the *adjacent* periods t and $t + 1$, the cumulated graph is affected at only one point—at period t . For periods prior to t , and for those later than (and including) $t + 1$, the cumulated utility level does not depend on the division of utility between t and $t + 1$. Hence, the evaluation of the redistribution in the additive case is modified only by the addition of one extra parameter:

$$\frac{\partial V}{\partial u_t} - \frac{\partial V}{\partial u_{t+1}} = w_t - w_{t+1} \pm \sigma_t < 0.$$

The coefficient σ_t is either added or subtracted depending on whether the average utility level in periods $t + 1$ through n is lower or greater than the level in the first t periods. If the future appears worse than the past from the vantage point of period t , then the addition of the (negatively valued) σ_t term creates an additional incentive for later scheduling of high-utility events.

3. *Separability with respect to outcome permutations.* The general model described here is consistent with a weaker form of separability, one that has an intuitive appeal in the context of scheduling decisions. The usual form of separability states that preference between two sequences that have elements in common does not depend on the nature of these common elements. The weaker version is expressed as follows: Suppose that two sequences of length n are permutations of the same set of events (i.e., utility levels) and that they also differ only over the first $m < n$ periods (or, equivalently, over the last m periods). Then, preference between them will not depend on the common ordering of events over the remaining $n - m$ periods.

Here is a concrete example of this type of *permutation separability*. A person is making plans for the next 2 months. At issue for the first month is how to order a fixed set of tasks that have to be completed. For the second month, a vacation month, the issue is how to order a fixed set of leisure activities. The condition we have just defined would imply that the optimal ordering of events for the second month is not affected by how the events are ordered in the first month (and vice versa).

It is not hard to demonstrate that our model generically satisfies this property. Because the two sequences under consideration are permutations of the same set of events, they will share the same reference line. Furthermore, sequences that agree on the scheduling for the first month will have the same decumulated utility graph for that segment. Regardless of the relation of this common segment to the reference line, it will contribute, through the deviation terms, exactly the same amount to the value of either such sequence and hence will have a zero net impact on preference among them. The optimal scheduling of events in the second segment is therefore independent of the schedule imposed in the initial one.

Two Studies of Preference for Outcome Sequences

In the following two studies, subjects rated the desirability of sequences defined by events occurring over five consecutive

weekends. The first study was designed to elicit basic properties of preferences toward sequences and to test the goodness of fit of the model presented in Equation 3. The second study was designed specifically to check for violations of preferential independence across periods.

Study 1

Method

Subjects. Fifty-two subjects recruited at the Museum of Science and Industry in Chicago were each given a ticket to a multimedia presentation (value \$4) in exchange for participating. Subjects represented a wide range of demographic characteristics. Forty-nine percent were male, 55% were married, and the age range was 18–58 years ($M = 32$). Six percent did not have a high school diploma, 11% had graduated from high school, 30% had some college but had not graduated, 36% had a college degree, and 17% possessed an advanced degree. The median yearly family income was approximately \$35,000.

Stimuli. Each sequence consisted of five consecutive weekends: one very enjoyable weekend, two moderately enjoyable weekends, and two boring weekends. All permutations of these five events created 30 distinct sequences. Subjects were told first to examine all 30 sequences and locate the best and worst one. The best was given a rating of 10 and the worst a 1. After this, subjects were instructed to go through the remaining 28 sequences and rate them according to their desirability relative to the best and worst. The exact instructions were as follows:

Imagine it is Monday and you are contemplating the next five weekends. Your situation is depicted in the diagram below.



Weekend: 1 2 3 4 5

The tall box represents a very pleasurable way to spend a weekend. Think of who you would like to spend such a weekend with and how you would like to spend it. In the diagram, you are doing the pleasurable thing on the third weekend.

The medium height box represents a moderately pleasurable way of spending a weekend. Think of something you enjoy, but do not love, doing. In the diagram, the first and fourth days are moderately pleasurable.

The flat line represents a boring and not at all fun way to spend a weekend. Think of something you do not enjoy doing. In the diagram, the second and fifth weekends are boring and not at all fun.

On the pages below you will find a series of 30 different sequences consisting of very pleasurable, pleasurable, and boring weekends. In each sequence there is one very pleasurable, two moderately pleasurable, and two boring weekends.

We want you to rate each sequence on a scale from 1 to 10, where 1 is the worst and 10 is the best. First, go through all the sequences and identify the one you like best and the one you like least. For the one you like best, enter a "10" on the blank line. For the one you like least, enter a "1." Then examine each sequence and assign it a rating from 1 to 10. If you like two sequences equally, you should give them the same rating.

The 30 sequences were presented in one of two opposite orderings determined by a random drawing.

Results

Because our sequences involved only three possible outcomes, the underlying utility assignment had only one free parameter, which we could identify with the position of the moderately enjoyable weekend on a 0–2 scale, with 0 and 2 being the values of the boring and very enjoyable weekends, respectively. The descriptions of the outcomes and the height of the vertical bars by which they were visually represented in the survey were intended to suggest a midpoint utility assignment for the moderately enjoyable weekend (i.e., scale value of 1).

We tested the model in Equation 3 by means of a maximum likelihood procedure, which returned for each subject estimates of the coefficients β and σ , as well as the value of the intermediate outcome, constrained to lie between 0 and 2. Positive values of β indicate people who prefer sequences that improve over time; negative values indicate people who discount the future. Negative values of σ designate people who prefer to spread good outcomes over time; people with positive σ values prefer to concentrate them at one point in the sequence.

The estimated value of the middle outcome was fairly close to 1 ($M = 0.81$), indicating that most subjects evaluated the moderately pleasurable weekend as about midway between the very pleasurable and boring weekends. Estimated values of β were positive for 40 of 52 subjects (77%), indicating a preference for improvement. Thirty-five of 52 subjects (67%) exhibited negative values of σ . Of the 52 subjects, 22 (42%) conformed exactly to the conjectured pattern: $\beta > -\sigma > 0$; This corresponds to the segment “dislikes declining, likes improving” in Figure 3. By contrast, we would have expected only 1 of 8, or 6.5 subjects, to reveal this pattern by chance. The mean value of β was .28, and the mean value of σ was $-.13$.

The mean correlation between actual and predicted ratings across all subjects was .60. Forcing the value of the middle-valued outcome exactly equal to 1 and rerunning the analysis lowered the mean correlation between actual and predicted ratings from .60 to .52 and had almost no impact on the other parameters. Therefore, all subsequent analyses were conducted with the value of the intermediate outcome set equal to 1.

To put the goodness-of-fit figures into perspective, we ran two additional analyses in an attempt to fit the standard discounted utility model,

$$\text{Rating} = c + \alpha \sum_{t=1}^5 \delta^t u_t, \tag{12}$$

to the subjects' ratings. The parameters c and α map the underlying utility scale into the 1–10 rating interval; the parameter δ measures whether a subject places more weight on early periods ($\delta < 1$) or on later ones ($\delta > 1$).

As expected, a large fraction (65%) of the subjects revealed “discount factors” (δ) greater than 1, implying greater weighting for outcomes in later periods. The mean value for δ was 1.23. The mean correlation between predicted and actual ratings was .38, and constraining the δ parameter to be less than or equal to 1 (i.e., ruling out negative time discounting) reduced this correlation to .08. Clearly, our model performs significantly better than the conventional discounted utility model, particularly when positive time discounting is assumed.

A comprehensive picture of intersubject variability is given

in Figure 5. Each subject is indicated by a point on the scatterplot, with the direction of each point relative to the origin determined by the relative values of β and σ , and the distance from the origin by the goodness of fit (for that subject). Specifically, the x - and y -coordinates in the scatterplot represent the normalized regression coefficients,

$$\beta' = r \frac{\beta}{\sqrt{\beta^2 + \sigma^2}}, \quad \sigma' = r \frac{\sigma}{\sqrt{\beta^2 + \sigma^2}},$$

where r is the correlation between predicted and actual sequence ratings. If the ratings of all subjects were fit perfectly by the model, then all the points in the central scatterplot would be distributed along the ellipse drawn in the figure that has a fixed radius of 1. Subjects located close to the x -axis care about improvement but are relatively indifferent to spread. Those located close to the y -axis care about spread but not about improvement.

Looking at the scatterplot, one notices a fairly large cluster of subjects in the lower right part of the panel, whose points come close to the ellipse that identifies a perfect fit. These subjects have positive coefficients for improvement ($\beta > 0$) and negative coefficients for σ , indicating a preference for spreading outcomes. One can find subjects who are motivated by impatience, but such subjects are rare, and their preferences are not well explained by the model. Only 2 subjects' data, located in the lower left quadrant, display impatience and a high goodness of fit.

In general, the model provides reasonable goodness of fit to the data, especially considering that it includes only two free parameters (not including the constant). Inevitably, however, it will fit certain individuals better than others, and certain types of sequences will be predicted more accurately than others. Table 2 presents an analysis of the mean and mean absolute

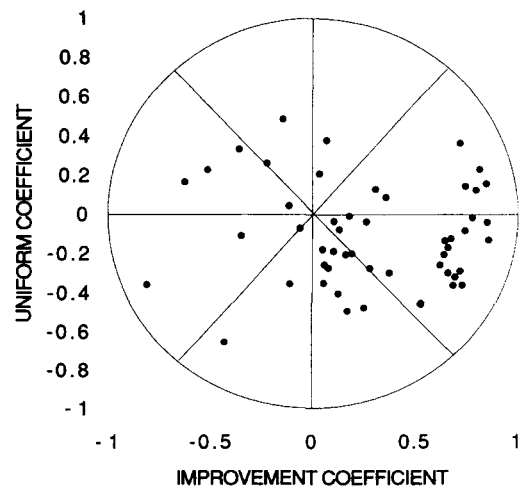


Figure 5. Scatterplot of normalized regression coefficients for all subjects in Study 1. (The x -axis plots the normalized coefficient for improvement and the y -axis the normalized coefficient for negative spread: $\beta' = r\beta/\sqrt{\beta^2 + \sigma^2}$, $\sigma' = r\sigma/\sqrt{\beta^2 + \sigma^2}$)

Table 2
Mean and Mean Absolute Regression Residuals by Sequence

Sequence	Mean rating	Mean predicted rating	Mean residual	Mean absolute residual	Mean squared residual
00112	6.11	6.14	-0.03	1.74	4.99
00121	5.74	6.11	-0.38	1.82	5.18
00211	5.57	6.08	-0.52	1.94	5.80
01012	6.81	6.11	0.70	1.69	5.20
01021	6.32	6.08	0.24	1.31	2.62
01102	6.08	6.08	-0.01	1.75	5.06
01120	5.36	5.63	-0.27	1.98	6.09
01201	5.68	5.73	-0.05	1.19	2.22
01210	5.58	5.31	0.28	1.74	4.97
02011	5.36	5.83	-0.47	1.42	3.52
02101	5.72	5.50	0.21	1.39	3.44
02110	5.55	5.08	0.47	1.78	4.53
10012	5.60	5.99	-0.38	1.45	3.27
10021	5.36	5.96	-0.60	1.49	3.98
10102	7.32	5.96	1.36	2.17	6.72
10120	5.51	5.50	0.00	1.50	3.68
10201	6.34	5.60	0.74	2.33	8.18
10210	5.32	5.18	0.14	1.34	3.14
11002	5.34	5.73	-0.39	1.87	5.72
11020	5.06	5.28	-0.22	1.64	4.35
11200	3.55	4.43	-0.88	1.69	4.30
12001	4.26	4.85	-0.60	1.48	3.49
12010	4.85	4.43	0.42	1.29	2.60
12100	3.62	3.91	-0.28	1.39	3.39
20011	4.62	5.18	-0.56	1.79	5.50
20101	5.13	4.85	0.28	1.66	4.48
20110	4.98	4.43	0.55	1.63	4.46
21001	4.11	4.33	-0.22	1.39	2.97
21010	4.43	3.91	0.53	1.43	3.53
21100	3.32	3.38	-0.06	1.29	2.91
<i>M</i>	5.29	5.29	0.0	1.62	4.34

regression residuals for the 30 sequences, averaged across all of the subjects. The mean error indicates the degree of bias in the model's prediction of mean ratings; the mean absolute and squared errors indicate the goodness of fit.

The pattern of residuals across sequences provides clues about aspects of preference that are not picked up by the model. Notably, the model overestimated the attractiveness of every sequence with two consecutive zero periods; among these, the sequence (1,1,2,0,0) was the most overvalued of any of the 30 patterns. This systematic error suggests that subjects may care about the relation between adjacent periods to a degree not captured by our global measure of spreading.

Study 2

Because the sequences in the first study were all permutations of a common set of outcomes, the results did not lend themselves to direct tests of separability. It was not possible to directly examine how the attractiveness of sequences changed when a single outcome was altered but all the other outcomes were held constant. The second study addressed this particular question.

In the study, subjects rated two matched blocks of sequences. The first, the *lean* block, consisted of a subset of 18 sequences from the previous study selected to represent the full diversity

of the original set of 30. The second, the *rich* block, contained a corresponding set of 18 sequences constructed by applying two operations: (a) inverting the utility of individual outcomes so that very enjoyable weekends occupied the position of the boring ones and vice versa and (b) flipping the order of weekends from last to first. Consequently, each sequence in the rich block contained one boring weekend, two intermediate weekends, and two highly enjoyable weekends. The specific reasons for this construction are explained now further.

For each block we created two orderings, one the reverse of the other. Labeling the two orderings *A* and *B* and the blocks *L* and *R*, each subject received one of the following four stimulus combinations: (*AL*, *BR*), (*AR*, *BL*), (*BL*, *AR*), and (*BR*, *AL*). Subjects were told to independently calibrate the rating scales for the two blocks, with 1 representing the worst and 10 the best sequence in each block.

Separability implies that a person who prefers, for example, the sequence (1,0,0,1,2) to (2,0,0,1,1) in the lean block should prefer the sequence (1,0,2,1,2) to (2,0,2,1,1) in the rich block because the common 0 in the third period has been replaced by the common 2. The two blocks of stimuli provided many equivalent pairs of this type.

The manner in which the two blocks of sequences were constructed also made possible a consistency check for our basic model (Equation 3). The two operations produced for each lean sequence a corresponding rich sequence that had the same amount of total improvement, as measured by $\sum d_i$, and the same total deviation from the reference line, as measured by $\sum |d_i|$. For example, starting with the lean sequence (1,0,0,1,2) utility inversion produces a rich sequence (1,2,2,1,0) that has the same deviation terms, d_i , but for a change in sign; formally, utility inversion is a linear transformation of the utility scale, $u(x) \rightarrow au(x) + b$, with $a = -1$, $b = +2$, which, as shown in the Appendix, has the effect of multiplying the deviation terms by the constant *a*. The second operation, order reversal, then restores the sign of the deviation terms but changes their order, from first to last. Neither operation affects the total absolute deviation of the sequence from its reference line. The effect of these transformations on the deviation terms is summarized as follows:

$$(d_1, d_2, d_3, d_4) \quad \sum d_i \quad \sum |d_i|$$

Original lean sequence:

$$(1,0,0,1,2) \quad (-.2, +.6, +1.4, +1.2) +3.0 \quad +3.4$$

Step 1:

Utility inversion:

$$(1,2,2,1,0) \quad (+.2, -.6, -1.4, -1.2) -3.0 \quad +3.4$$

Step 2:

Order reversal:

$$(0,1,2,2,1) \quad (+1.2, +1.4, +.6, -.2) +3.0 \quad +3.4.$$

According to the simple model, $\sum u_i + \beta \sum d_i + \alpha \sum |d_i|$, the value of a rich sequence equals the value of its corresponding lean sequence, plus an additive constant, which is the same for all sequences and represents the increase of total utility by two units in the rich set. After recalibration of the rating scale, the

nominal ratings assigned to each pair of corresponding lean and rich sequences should therefore be the same.

Method

Fifty-seven subjects were recruited at the Museum of Science and Industry in Chicago and were given a ticket to a multimedia presentation (value \$4) in exchange for participating. Forty-nine percent of the subjects were male, 55% were married, and the mean age was 33. Four percent did not have a high school diploma, 6% had graduated from high school, 33% had some college, 35% had a college degree, and 20% had earned advanced degrees. The median yearly family income of the sample was again approximately \$35,000.

Results

The success of the model and the distribution of parameter estimates were comparable to those obtained in Study 1. Again, most subjects preferred improving sequences (81% and 79% for the lean and rich sets, respectively). The mean goodness of fit, as measured by the correlation between actual and expected values, was .61 for the lean sequence set and .58 for the rich set (i.e., slightly higher than that observed in the first experiment). Figure 6 displays the β and σ values and the goodness of fit for individual subjects in the same format as in Figure 5. The dots represent estimates derived from the lean block of sequences; the plus signs result from the rich block. The correlation between the β coefficient values estimated from the two blocks of data was .64; the correlation for the σ coefficients was .44.

Table 3 presents the mean ranks and ratings for the 18 pairs of corresponding sequences in the two blocks. The results are largely consistent with the simple model. The two best sequences, (0,1,1,0,2) and (0,2,1,1,2), are a corresponding pair, as are the two worst sequences, (2,1,1,0,0) and (2,2,1,1,0). The correlation across corresponding pairs in the two blocks is .95 (for mean ranks) and .92 (for mean ratings). The only two cases in which mean ranks (or ratings) diverge by more than 1.0

involve the two lean sequences (0,1,1,2,0) and (0,2,1,1,0), which both start and end on a bad note. Their corresponding sequences, (2,0,1,1,2) and (2,1,1,0,2), have good outcomes in these salient positions, which may enhance their value.

It is worth noting that the predicted one-to-one relation between the sequences in the lean and rich blocks is not self-evident, outside of the context of our model. For many pairs in the list, for example, (1,1,0,0,2) and (0,2,2,1,1) or (0,2,0,1,1) and (1,1,2,0,2), it is not at all obvious why they should receive similar ratings within their respective blocks. Consequently, the fact that the ratings *do* agree for so many of the pairs is strong evidence in support of the simple model.

Turning now to the separability tests, we examined all pairs of sequences that differed in exactly one period. Given the construction of the lean and rich blocks, each sequence in such a pair would have either a 0 or a 2 in the nonmatching position. It is natural to further divide these pairs into five subgroups, depending on whether the nonmatching weekend occupied the first, second, third, fourth, or fifth position. The ranks were computed within each of these subdivisions and separately for the lean and the rich blocks.

Separability implies that such matching pairs of sequences should have the same rank within their respective blocks. Violations of independence are thus revealed by a divergence in ranking. This is a conservative test of independence violations because preferential interactions might occur but would not be sufficiently strong to influence ranks.

Table 4 shows the average sequence rankings in a format relevant to our separability tests. (We converted the direct ratings into rankings in an obvious manner, giving sequences with the same rating a tie rank.) The first pair of columns present mean ranks for all sequences that were identical except for the first period, which was either boring (first column) or very enjoyable (second column). The numbers 5.5 and 5.9, for example, represent the mean ranks of sequences (0,0,1,1,2) and (2,0,1,1,2). The mean value of each of these columns is the mean of the ranks 1–8. Likewise, the second pair of columns provides mean ranks for sequences that are identical except in the second period. If independence is satisfied, then the rank of the matched pairs should not be significantly different from one another. The underlined element in each pair is the one that has lower total deviation from its respective reference line (i.e., $\sum |d_i|$), and therefore it is also the one that should have higher rank according to our theory.

The data in Table 4 provide strong evidence of preferential interactions of the type predicted by the model. For the 36 pairs of ranks, in 29 cases (or 81%) the difference in ranks has the correct sign, including 13 of the 14 pairs where the difference is significant at the .05 level.

For example, introducing a very pleasurable weekend in the first position of a sequence in which the second weekend is boring, in both cases improves its relative ranking, whereas providing such a weekend when the first weekend is already very enjoyable, lowers its relative ranking. In the second pair of columns, substitution of the high-value outcome in the second period helps the sequence (1,*0,1,2), which is otherwise very thin in the first four periods, and hurts the sequence (2,*1,1,0), which adopts a monotonically declining profile when the very enjoyable weekend is inserted. Perhaps the most




Figure 6. Scatterplot of normalized regression coefficients for all subjects in Study 2.

Table 3
Mean Ranks and Ratings for Corresponding Pairs of Sequences in the Lean and Rich Blocks

Lean sequence	Corresponding rich sequence	Mean ranks		Mean ratings	
		Lean	Rich	Lean	Rich
01102	02112	13.3	12.5	7.2	7.1
00112	01122	13.1	11.6	7.2	6.3
01201	12012	12.3	12.4	6.6	6.9
01021	10212	11.7	12.2	6.3	7.0
00211	11022	11.3	11.0	5.9	6.3
02011	11202	10.9	10.1	5.9	5.8
10210	21021	10.3	10.2	5.6	6.1
10012	01221	10.1	10.0	5.5	5.9
11020	20211	10.1	9.1	5.5	5.3
01120	20112	9.7	11.4	5.2	6.4
11002	02211	9.0	8.9	5.0	5.2
12010	21201	8.6	9.2	4.7	5.4
02110	21102	8.5	9.8	4.6	5.8
20110	21120	8.1	7.7	4.6	4.4
20011	11220	6.9	6.5	3.8	3.9
21001	12210	6.5	6.9	3.6	4.1
11200	22011	6.1	6.9	3.4	4.2
21100	22110	4.5	4.6	2.7	2.8

clear evidence of the importance of spread comes from the third pair of columns, where the middle period is being changed. Independence is violated at the .05 level for six of the eight residual subsequences. Likewise, in the fourth column, there are several sequences that show large interaction effects. For example, the sequence (2,1,1,*0) becomes maximally declining and poorly spread on the introduction of a boring weekend in the fourth period. The two back-loaded sequences (0,1,1,*2) and (0,1,2,*1), however, are both hurt by the introduction of the enjoyable weekend in the fourth position. Finally, in the fifth pair of columns, there are two dramatic interactions, (1,1,0,2,*), which is hurt by the introduction of a very enjoyable weekend in the last period, and (2,1,1,0,*), which is helped.

For a more global test of interaction, we ran five ANOVAs, one for each of the columns summarized in Table 3. The value of the variable period (either 0 or 2) was the first variable, and the pattern of the residual sequence was the second. We also included subject identification as an independent variable, to adjust for differences in subjects' mean ratings.

If independence holds, we would expect to observe main effects for the value of consumption in the first period and for the remaining sequence but should not observe a significant interaction between the two. In other words, a subject's ranking of a sequence should depend on what happens in the first period and what happens in subsequent periods but should not depend on an interaction between the first period and subsequent periods.

The independence hypothesis was, in fact, rejected in every period except the first in which the interaction term bordered on significance, $F(7, 705) = 1.60, p < .14$. The interaction effect was strongly significant in the second, $F(5, 517) = 5.34, p < .0001$, third, $F(7, 705) = 6.92, p < .0001$, fourth, $F(5, 517) = 7.00, p < .0001$, and fifth, $F(7, 705) = 3.2, p < .003$, periods.

Discussion

Several general conclusions emerge from these analyses. It appears that for sequences of outcomes, negative time preference is the rule rather than the exception. In the two studies just presented, more than 75% of the subjects exhibited a desire for improvement. We also observed strong violations of temporal independence, which indicates that any model that excludes interactions, such as an additive utility model, will miss important aspects of preferences.

Finally, we found that our two-parameter model did a good job of explaining individual's preferences toward sequences of outcomes. It did better than the discounting formula, even if the discount rate was allowed to be negative. As shown in the second study, it accurately predicted which sequences would receive the same rank in separate choice sets.

Table 4
Test of Preferential Independence Mean Rank by Sequence

Residual sequence	Value of first period		Value of second period		Value of third period		Value of fourth period		Value of fifth period	
	0	2	0	2	0	2	0	2	0	2
	0112	5.5	5.9	4.6	4.7			5.0	4.1**	5.3
0121					5.7	4.8***	4.7	3.6***		
0211	4.4	4.9	3.9	3.5*	5.3	4.3**			4.7	5.1
1012			3.4	4.6***	5.0	6.0**				
1021	4.7	5.4**							5.5	4.9
1102	5.6	5.0*			4.4	5.0	3.6	4.1*	5.5	4.1***
1120	3.7	3.9			5.0	3.3***	2.8	2.5	3.4	3.8
1201	4.8	4.7							4.7	4.9
1210			3.6	3.0**	4.1	3.5				
2011	4.2	3.7*	2.5	3.1**	3.2	4.4***			4.3	4.4
2101					3.1	4.5***	3.0	3.7***		
2110	3.0	2.4	2.9	2.2**			2.1	3.0***	2.5	3.8***

Note. The boldfaced element in each pair should have higher rank according to theory.
* $p < .10$. ** $p < .05$. *** $p < .01$.

At the same time, however, we acknowledge several limitations in our work. First, there are subjective considerations that may be important to individuals but that are excluded from the model. For example, the fact that our model tended to predict overly high ratings of sequences with two successive boring periods suggests that our spread measure was not sufficiently sensitive to such interactions between adjacent periods. The underprediction of the sequence (1,0,2,0,1)—compare Table 2—may indicate a premium for symmetry and so forth.

Second, our studies involved outcomes that were all gains; it is an open question how well the model generalizes to sequences involving losses or combinations of gains and losses. It is possible, for example, that people will prefer to isolate a single gain in a sequence of losses or a single loss in a sequence of gains, a pattern that would not be consistent with our model.

Third, our sequences were all permutations of a common set of outcomes over a constant number of periods of equal duration. Varey and Kahneman (in press) argued that preferences tend to be insensitive to the duration of a sequence and observed violations of dominance that stem, in part, from this insensitivity. This important finding needs to be replicated and, if robust, addressed theoretically.

Fourth, we did not address whether individuals would actually stick to the choice of an improving utility sequence when opportunities to change their original choice appeared along the way. The choice of an improving sequence represents a decision to defer gratification, and we know from research and from personal experience that such decisions are easier to entertain than to carry out. Countervailing temptations arise partly from the shape of the psychological discount function, which places disproportionate weight on immediate satisfaction (Ainslie, 1975), and partly from other factors, such as sensory contact with rewards (Mischel, 1974). Mischel drew a useful distinction between *delay choice*—choice between immediate and delayed alternatives—and *delay of gratification*, which refers to the implementation of such delay choices. He showed that implementing an initial decision to delay is not conflict free if the immediate reward remains available during the waiting period. Consequently, although individuals may prefer sequences that improve, in reality they may succumb to the manifold temptations of the moment. Repeated lapses of this type will, of course, in the long run produce a declining utility sequence, precisely the opposite of that which is desired.

However, individuals are not helpless in the face of their urges. They may resort to self-control strategies (Ainslie, 1975; Schelling, 1984), cognitive restructuring (Ainslie, 1975; Mischel, 1974), and raw willpower (Hoch & Loewenstein, 1991). Many social institutions and arrangements provide further support for delay of gratification. These include social security, retirement plans, term insurance, Christmas and dieting clubs (Thaler & Shefrin, 1981), and traditionally increasing wage trajectories (Frank, in press; Loewenstein & Sicherman, 1991). All of these combine to help people avoid declining utility profiles.

Conclusion

Understanding choice between sequences is important because planning for the future invariably requires one to choose

between alternative sequences of outcomes. Taking a vacation now may forestall a future vacation; increasing one's spending in the present may force reductions in future expenditures; and dieting in the present is intended to produce delayed rewards, whereas binge eating entails delayed costs. In each of these cases, a given decision has multiple consequences that are spread out over time.

Previous empirical work on time preference has focused almost entirely on the tradeoff that arises when two outcomes of different values and occurring at different times are compared. The tacit premise has been that such judgments will reveal an individual's "raw" time preference, from which one can then derive preferences over more complex objects (e.g., retirement plans, intertemporal income profiles, etc.). This view we now know is fundamentally incorrect. The empirical evidence presented in this article, in conjunction with the related work of Frank (in press), Loewenstein and Sicherman (1991), Ross and Simonson (1991), and Varey and Kahneman (in press), shows that as soon as an intertemporal tradeoff is embedded in the context of two alternative *sequences* of outcomes, the psychological perspective, or "frame," shifts, and individuals become more farsighted, often wishing to postpone the better outcome until the end. The same person who prefers a good dinner sooner rather than later, if given a choice between two explicitly formulated sequences, one consisting of a good dinner *followed* by an indifferent one, the other of the indifferent dinner *followed* by the good one, may well prefer the latter alternative. Sequences of outcomes that decline in value are greatly disliked, indicating a negative rate of time preference.

The sensitivity of intertemporal decisions to choice representation has important policy implications. Efforts to lengthen time perspectives have generally focused on material inducements. For example, attempts to increase the personal savings of Americans have typically involved tax deductions on certain types of interest income. Although not denigrating such efforts, our research suggests that there may be other, more effective and less costly methods of altering time perspective. Such methods can take the form of media and educational campaigns that express decisions as sequences rather than as individual decisions. Alternatively, policymakers could bolster and expand on the institutions that already exist that implicitly or explicitly present decision makers with choices between sequences.

The significant difference in preferences observed in intertemporal choices involving single outcomes and sequences challenges the claim often made that groups and individuals differ in their fundamental attitude toward the future. It is plausible that such variations are not attributable to any fundamental attitude toward the future but instead reflect differences in the way that options are perceived. Any factor, whether personal or situational, that causes intertemporal choices to fragment and to be perceived as a series of individual decisions will tend to induce high positive time discounting. Likewise, factors that cause such decisions to be internally "framed" as sequences will promote low and even negative time discounting.

References

- Ainslie, G. (1975). Specious reward: A behavioral theory of impulsiveness and impulse control. *Psychological Bulletin*, 82, 463–509.

- Barnes, O., & Barnes, L. W. (1964). Choice of delay of inevitable shock. *Journal of Abnormal Social Psychology*, 68, 669–672.
- Bell, D. E. (1977). A utility function for time streams having inter-period dependencies. *Operations Research*, 25, 448–458.
- Bordley, R. F. (1986). Satiation and habit persistence (or the dieter's dilemma). *Journal of Economic Theory*, 38, 178–184.
- Carlsmith, J. M. (1962). *Strength of expectancy: Its determinants and effects*. Unpublished doctoral dissertation, Harvard University, Cambridge, MA.
- Carson, R. T., Horowitz, J. K., & Machina, M. J. (1987). *Discounting mortality risks*. Unpublished manuscript.
- Elster, J. (1985). Weakness of the will and the free-rider problem. *Economics and Philosophy*, 1, 231–265.
- Elster, J., & Loewenstein, G. (in press). Utility from memory and anticipation. In G. Loewenstein & J. Elster (Eds.), *Choice over time*. New York: Russell Sage Foundation.
- Epstein, L. G., & Hynes, J. A. (1983). The rate of time preference and dynamic economic analysis. *Journal of Political Economy*, 91, 611–635.
- Frank, R. (1989). Frames of reference and the quality of life. *American Economic Review*, 79, 80–85.
- Frank, R. (in press). Frames of reference and the intertemporal wage profile. In G. Loewenstein & J. Elster (Eds.), *Choice over time*. New York: Russell Sage Foundation.
- Gilboa, I. (1989). Expectation and variation in multi-period decisions. *Econometrica*, 57, 1153–1169.
- Helson, H. (1964). *Adaptation-level theory: An experimental and systematic approach to behavior*. New York: Harper & Row.
- Hoch, S., & Loewenstein, G. (1991). Time-inconsistent preferences and consumer self-control. *Journal of Consumer Research*, 17, 1–16.
- Horowitz, J. K. (1988). *Discounting money payoffs: An experimental analysis*. Unpublished manuscript.
- Hsee, C. K., & Abelson, R. P. (1991). The velocity relation: Satisfaction as a function of the first derivative of outcome over time. *Journal of Personality and Social Psychology*, 60, 341–347.
- Hsee, C. K., Abelson, R. P., & Salovey, P. (1991). The relative weighting of position and velocity in satisfaction. *Psychological Science*, 2, 263–266.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 363–391.
- Koopmans, T. C. (1960). Stationary ordinal utility and impatience. *Econometrica*, 28, 287–309.
- Koopmans, T. C., Diamond, P. A., & Williamson, R. E. (1964). Stationary utility and time perspective. *Econometrica*, 46, 82–100.
- Linville, P., & Fischer, G. (1991). Preferences for combining or separating events: A social application of prospect theory and mental accounting. *Journal of Personality and Social Psychology*, 60, 5–23.
- Loewenstein, G. (1987). Anticipation and the valuation of delayed consumption. *Economic Journal*, 97, 666–684.
- Loewenstein, G., & Sicherman, N. (1991). Do workers prefer increasing wage profiles? *Journal of Labor Economics*, 9, 67–84.
- Meyer, R. F. (1976). Preferences over time. In R. L. Keeney & H. Raiffa (Eds.), *Decisions with multiple objectives: Preferences and value trade-offs* (pp. 473–485). New York: Wiley.
- Meyer, R. F. (1977). State dependent time preference. In D. Bell, R. L. Keeney, & H. Raiffa (Eds.), *Conflicting objectives in decision* (pp. 232–244). New York: Wiley.
- Miller, N., & Campbell, D. T. (1959). Recency and primacy in persuasion as a function of the timing of speeches and measurements. *Journal of Abnormal and Social Psychology*, 59, 1–9.
- Mischel, W. (1974). Processes in delay of gratification. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 7, pp. 249–292). San Diego, CA: Academic Press.
- Prelec, D., & Loewenstein, G. (1991). Decision-making over time and under uncertainty: A common approach. *Management Science*, 37, 770–786.
- Ross, W. T., Jr., & Simonson, I. (1991). Evaluations of pairs of experiences: A preference for happy endings. *Journal of Behavioral Decision Making*, 4, 273–282.
- Samuelson, P. (1937). A note on measurement of utility. *Review of Economic Studies*, 4, 155–161.
- Schelling, T. (1984). Self-command in practice, in policy, and in a theory of rational choice. *American Economic Review*, 74, 1–11.
- Thaler, R., & Johnson, E. (1990). Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice. *Management Science*, 36, 643–660.
- Thaler, R., & Shefrin, H. (1981). An economic theory of self-control. *Journal of Political Economy*, 89, 392–410.
- Tversky, A., & Griffin, D. (1991). Endowment and contrast in judgments of well-being. In R. J. Zeckhauser (Ed.), *Strategy and choice* (pp. 297–318). Cambridge, MA: MIT Press.
- Varey, C., & Kahneman, D. (in press). The integration of aversive experiences over time: Normative considerations and lay intuitions. *Journal of Behavioral Decision Making*.

(Appendix follows on next page)

Appendix

To show that the two forms of the model are equivalent, we start with Equation 6 and substitute for d_i^+ and d_i^- (from Equation 4):

$$\begin{aligned} V(X) &= \sum_{i=1}^n u_i + \sum_{i=1}^{n-1} \sigma_i^+ d_i^+ + \sum_{i=1}^{n-1} \sigma_i^- d_i^- \\ &= \sum_{i=1}^n u_i + \sum_{i=1}^{n-1} \left(\frac{\sigma_i^+ - \sigma_i^-}{2} \right) d_i + \sum_{i=1}^{n-1} \left(\frac{\sigma_i^+ + \sigma_i^-}{2} \right) |d_i|. \end{aligned}$$

Substituting for d_i and σ_i ,

$$d_i = \frac{t}{n} \sum_{i=1}^n u_i - \sum_{i=1}^t u_i, \quad \sigma_i = \frac{\sigma_i^+ + \sigma_i^-}{2},$$

yields the second form of the model (Equation 7),

$$\begin{aligned} V(X) &= \sum_{i=1}^n u_i + \sum_{i=1}^{n-1} \left(\frac{\sigma_i^+ - \sigma_i^-}{2} \right) \left(\frac{t}{n} \sum_{i=1}^n u_i - \sum_{i=1}^t u_i \right) + \sum_{i=1}^{n-1} \sigma_i |d_i| \\ &= \sum_{i=1}^n w_i u_i + \sum_{i=1}^{n-1} \sigma_i |d_i|, \end{aligned}$$

with the weights equal to

$$w_i = 1 + \sum_{i=1}^{n-1} \frac{i}{n} \left(\frac{\sigma_i^+ - \sigma_i^-}{2} \right) + \sum_{i=t}^{n-1} \left(\frac{\sigma_i^+ - \sigma_i^-}{2} \right).$$

The first two terms do not depend on the period t . Therefore, the difference between adjacent weights simplifies to

$$\begin{aligned} w_t - w_{t+1} &= - \sum_{i=t}^{n-1} \left(\frac{\sigma_i^+ - \sigma_i^-}{2} \right) + \sum_{i=t+1}^{n-1} \left(\frac{\sigma_i^+ - \sigma_i^-}{2} \right) \\ &= \sigma_t^+ - \sigma_t^- \end{aligned}$$

To check that $u(x)$ is an interval scale, we first assess the impact of the substitution of $u_i^*(x_i) = au(x_i) + b$ for $u(x_i)$ on the following deviation terms:

$$\begin{aligned} d_i^* &= \frac{t}{n} \sum_{i=1}^n (au_i + b) - \sum_{i=1}^t (au_i + b) \\ &= a \frac{t}{n} \sum_{i=1}^n u_i + tb - a \sum_{i=1}^t u_i - tb \\ &= a \left(\frac{t}{n} \sum_{i=1}^n u_i - \sum_{i=1}^t u_i \right) \\ &= ad_i. \end{aligned}$$

After the transformation, the value function (Equation 7) appears as

$$\begin{aligned} V^*(X) &= \sum_{i=1}^n w_i (au_i + b) + \sum_{i=1}^{n-1} \sigma_i a |d_i| \\ &= a \left(\sum_{i=1}^n w_i u_i + \sum_{i=1}^{n-1} \sigma_i |d_i| \right) + b \sum_{i=1}^n w_i \\ &= aV(X) + b \sum_{i=1}^n w_i, \end{aligned}$$

which ranks sequences in the same way as does $V(X)$.

Received August 5, 1991

Revision received October 21, 1991

Accepted November 4, 1991 ■