Anomalies
Intertemporal Choice

George Loewenstein and Richard H. Thaler

Economics can be distinguished from other social sciences by the belief that most (all?) behavior can be explained by assuming that agents have stable, well-defined preferences and make rational choices consistent with those preferences in markets that (eventually) clear. An empirical result qualifies as an anomaly if it is difficult to "rationalize," or if implausible assumptions are necessary to explain it within the paradigm. This column will present a series of such anomalies. Readers are invited to suggest topics for future columns by sending a note with some references to (or better yet copies of) the relevant research. Comments on anomalies printed here are also welcome. The address is: Richard Thaler, c/o *Journal of Economic Perspectives*, Johnson Graduate School of Management, Malott Hall, Cornell University, Ithaca, NY 14853.

Introduction

Intertemporal choices, decisions in which the timing of costs and benefits are spread out over time, are both common and important. How much schooling to obtain, whom to marry, whether to have children, how much to save for retirement, how to invest, whether to buy a house, and if so which house to buy—all these vital decisions have strong intertemporal components. As examples of individual decision

■ *George Loewenstein is Associate Professor, Graduate School of Business, University of Chicago, Chicago, Illinois and Visiting Scholar, Russell Sage Foundation, New York, New York. Richard Thaler is Henrietta Johnson Louis Professor of Economics, Johnson Graduate School of Management, Cornell University, Ithaca, New York.*

making, intertemporal choices are also interesting because the relevant economic theory makes unusually testable predictions. In many contexts, economic theories of individual behavior are untestable because the predictions are too vague. Almost any choice, no matter how bizarre, can be rationalized by finding some utility function for which the choice represents an optimal solution. In contrast, for decisions involving choices between time streams of money (receipts and payments), economic theory makes a precise and testable prediction, namely that (at the margin) people should discount money streams at the (after-tax) market rate of interest (r).

The existence of capital markets creates what amounts to an internal arbitrage opportunity for the consumer. If presented with an investment option that pays off at a rate higher than r, the consumer can enjoy greater consumption in every period by accepting the option and borrowing appropriately at rate r. Options that pay less than r should be rejected since they are dominated by lending in the capital market. The implication is that consumers should make intertemporal trade-offs so that their marginal rate of time preference equals the interest rate. Furthermore, consumers should be consistent in their intertemporal choices. The discount rate used should be constant across situations and over time. However, research shows that depending on the context examined, the implied discount rates of observed behavior can vary from negative to several hundred percent per year.

A well-known example of apparent negative discount rates is the fact that a large majority of U.S. taxpayers receive refunds every year from the Internal Revenue Service. These interest-free loans to the government are easily avoidable by adjusting the withholding rate. Similarly, many school teachers are given the choice between being paid in 9 monthly installments (September-June) or 12 (September-August). Most of those given this choice elect the latter option. Finally, studies of life-cycle consumption choices reveal that consumption tends to increase over time until retirement. In the absence of binding borrowing constraints, this pattern can only be consistent with the life-cycle theory if people have negative discount rates (see, Courant, Gramlich, and Laitner, 1986).

Examples of extremely high discount rates are also easy to find. A recent change in West Virginia law provides an example. Students under the age of 18 who drop out of school lose their driving permits. The first year results indicate that this law has reduced the dropout rate by one-third. It seems implausible that one-third of the high school dropouts were so close to the margin that the loss of driving privileges for a year or two (or more precisely, the expected costs of driving illegally for this period) could tip a rational human capital investment decision toward completing high school. Rather, the behavior seems to reveal extremely myopic preferences. A similar myopia is evident in the lament of a dermatologist that her warnings about the risk of skin cancer have little effect, but "My patients are much more compliant about avoiding the sun when I tell them that it can cause large pores and blackheads."

It is not just teenagers and sun lovers who display high discount rates. Most homeowners have too little insulation in their attics and walls, and fail to buy more expensive energy-efficient appliances even when the pay-back period for the extra expense is less than a year. Hausman's (1979) study of air conditioner purchases,

which examined consumer tradeoffs between purchase price and delayed energy payments, estimated an average consumer discount rate of about 25 percent. A subsequent study by Gately (1980) comparing pairs of refrigerators differing only in energy use and initial purchase price revealed that the implicit discount rates associated with purchasing the cheaper models were incredibly high: from 45 to 130 percent assuming an electricity cost of 3.8 cents per kilowatt hour, and from 120 to 300 percent at 10 cents per kilowatt hour. Most recently, Ruderman, Levine and McMahan (1986) computed the discount rates implicit in several different kinds of appliances (for the average model on the market, relative to the most efficient): space heaters, air conditioners, water heaters, refrigerators and freezers. They found that the implicit discount rate for room air conditioners was 17 percent, somewhat lower than Hausman's estimate. However, the discount rates for other appliances were much higher, e.g., gas water heater, 102 percent; electric water heater, 243 percent; and freezer, 138 percent. Economic theory has a clear prediction about these inefficient appliances—they will not be produced. But they are produced, and purchased.[1]

So, as usual, where there are testable predictions, there are anomalies. The remainder of this column examines a number of situations in which people do not appear to discount money flows at the market rate of interest or any other single discount rate. Discount rates observed in both laboratory and field decision-making environments are shown to depend on the magnitude and sign of what is being discounted, on the time delay, on whether the choice is cast in terms of speed-up or delay, on the way in which a choice is framed, and on whether future benefits or costs induce savoring or dread.

Variations in the Discount Rate for an Individual

An experiment that investigated the first three of these effects was presented in Thaler (1981). Subjects (mostly students) were asked to imagine that they had won some money in a lottery conducted by their bank. They could take the money now or wait until later. They were asked how much they would need to be paid to make waiting as attractive as immediate payment. Each subject received a 3 × 3 table to fill in with amounts of money varied along one dimension and length of time along the other. Four versions of the questionnaires were used, three involving gains, and one involving losses. In the losses version, subjects were asked to imagine that they had been issued a traffic fine that could either be paid at face value now or at an increased

---

[1]Two other explanations might be offered for the purchase of inefficient appliances: ignorance and illiquidity. According to the ignorance hypothesis, customers do not know, or bother to find out, the advantages of buying a more efficient model even though that information is plainly displayed on government mandated labels. According to the illiquidity argument, customers are so short of cash that they cannot afford to buy the more efficient model. (Of course, these are precisely the customers who cannot afford to buy the cheaper model!) Since most appliances are probably purchased on credit, and since the extra cost of the energy efficient model is relatively small, it seems unlikely that borrowing constraints are really the answer.

price later. In all cases subjects were asked to assume that there was no risk of not getting the reward (or of avoiding the fine) if they waited. All amounts were to be received (or paid) by mail.[2] The experiment thus manipulated the three variables of interest: the length of time to be waited; the magnitude of the outcome; whether the outcome is a gain or loss.

Three strong patterns emerged from the subjects' responses. First, discount rates declined sharply with the length of time to be waited, consistent with earlier findings for animals (Herrnstein, 1961; Ainslie, 1975). Second, discount rates declined with the size of the reward. Discount rates for small amounts (under $100) were very high, while those for larger amounts were more reasonable. Third, discount rates for gains were much higher than for losses. Subjects needed to be paid a lot to wait for a reward, but were unwilling to pay very much to delay a fine.

These three findings have been replicated in a much larger study by Benzion, Rapoport, and Yagil (1989). They used a 4 × 4 × 4 design which manipulated the time delay (0.5, 1, 2, and 4 years), amount of money ($40, 200, 1000, and $5000), and scenario (postponing a gain; postponing a loss; expediting a gain; and expediting a loss). The subjects were undergraduate and graduate students in economics and finance at two Israeli universities, a relatively sophisticated subject pool. Their results are shown in Figure 1 (averaging across the four scenarios). As can be seen clearly, discount rates again decline sharply with the length of time to be waited and the size of the prize.[3]

We will discuss each of these three strong patterns of discount rate variations in turn.

Dynamic Inconsistency

The negative relationship between discount rates and time delay has important consequences for the dynamic consistency of behavior. Suppose, as illustrated in Figure 2, that an individual must choose between two rewards, a small early reward S,

---

[2] In this study, and some others described here, the questions asked were hypothetical. Of course, all things being equal it would be better to study actual choices. However, there are serious trade-offs between hypothetical and real money methods. Using hypothetical questions one can ask subjects to consider options that incorporate large amounts of money, both gains and losses, and delays of a year or more. In studies using real choices, the experimenter must reduce the size of the stakes and the length of the delay, and it is difficult to investigate actual losses. Also, in a hypothetical question, one can ask the subject to assume that there is no risk associated with future payments, while in experiments using real stakes, subjects must assess the experimenter's credibility. It is reassuring that in this domain, as well as many others, the phenomena discovered using hypothetical choices have been reproduced in studies using actual choices, see for example, Horowitz (1988), and Holcomb and Nelson (1989).

[3] It is obvious that whatever pattern of choices subjects indicate in these experiments, market interest rates do not depend (greatly) on either magnitude or time delay, but this does not imply that the experimental evidence is irrelevant for economics. Economics is concerned with predicting both market prices and individual behavior. Though arbitrageurs may assure that one cannot earn (much) more interest from buying and selling a series of 12 one-month treasury bills than a single one-year bond, this does not guarantee that predictions at the individual level will be accurate. If car customers elect financing over more attractive rebates, no (costless) arbitrage opportunity exists for anyone else. A bank could try to convince car buyers that they would be better off taking the rebate and financing the purchase at their bank, but such campaigns are expensive, and consumers may be skeptical regarding the impartiality of the advice they are being given.

Figure 1

**Discounting as a Function of Time Delay and Money Amount.**
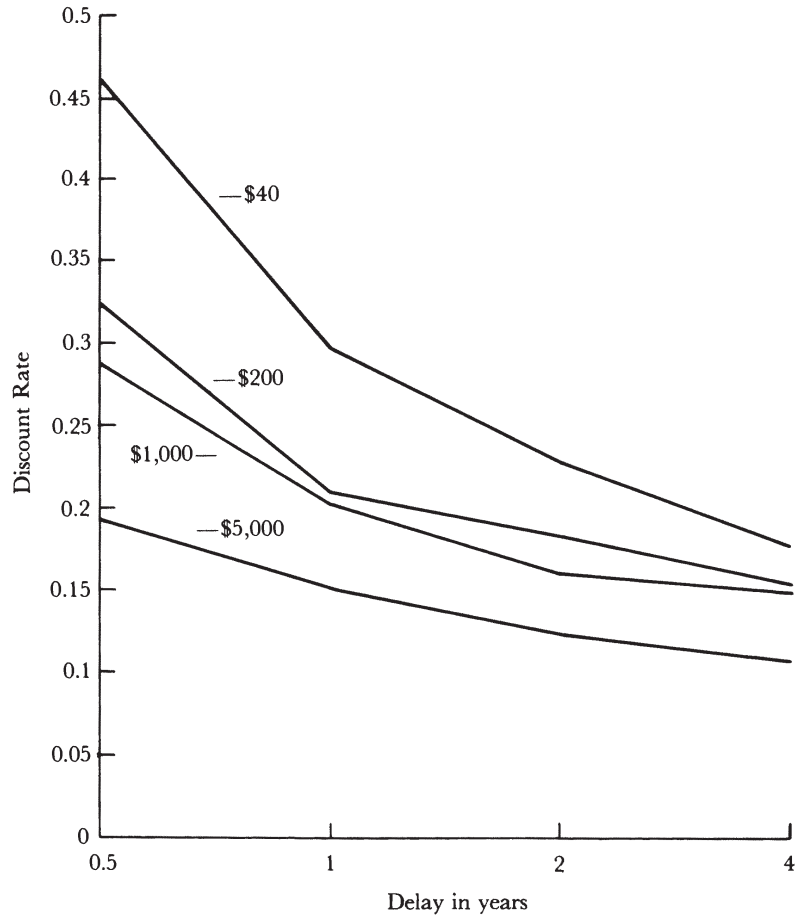
Source: Benzion et al. (1989).

which occurs at t_1, and a bigger later reward B, which occurs at t_2.[4] The lines represent the present utility of the rewards as perceived by the individual at different points in time. If the individual discounts the future at a constant rate, that is, if discounting is constant for different time delays, then the curves will never cross. However, if discounting decreases as a function of time delay, as the empirical research suggests, then the curves may cross, leading to a reversal of preference. When both rewards are sufficiently distant, the individual prefers B, but as S becomes more proximate, its relative value increases until at t^*, S abruptly comes to dominate B in terms of present utility. The significance of the crossing curves is that behavior will not generally be consistent over time. In the morning, when temptation is remote, we vow to go to bed early, stick to our diet, and not have too much to drink. That night we stay out until 3:00 a.m., have two helpings of chocolate decadence, and sample every variety of Aquavit at a Norwegian restaurant. Applied to saving, as Strotz (1956)

[4]This analysis is based on Ainslie (1975).

Figure 2
Non-Exponential Discounting.
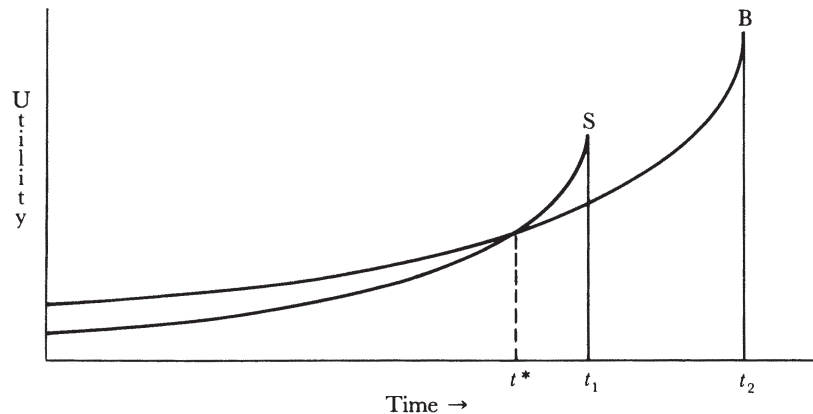
Source: Ainslie (1975).

demonstrated, if the discount rate declines over time, then people will always consume more in the present than called for by their previous plans.

The problem of dynamic inconsistency raises questions about consumer sovereignty. Who is sovereign, the self who sets the alarm clock to rise early, or the self who shuts it off the next morning and goes back to sleep? It is instructive that we normally see the far-sighted self take actions which constrain or alter the behavior of the myopic self. Dieters pay money to stay on "fat farms" whose main appeal is that they guarantee to underfeed their guests; alcoholics take antabuse which causes nausea and vomiting if they take a drink; smokers buy cigarettes by the pack (rather than by the carton which is cheaper). And, though no longer fashionable, for many years Christmas clubs were extremely popular in the U.S. These savings plans offered the unusual combination of inconvenience (deposits were made in person every week), illiquidity (funds could not be withdrawn until late November), and low interest (in some cases, zero interest). Of course, illiquidity was the Christmas club's raison d'être since customers wanted to assure themselves of funds to pay for Christmas presents. Recognizing the limited ability of conventional decision models to account for self-binding behavior and other forms of intrapersonal conflict, a number of authors have proposed models that view economic behavior as an internal struggle between multiple selves with conflicting preferences (Ainslie 1975, forthcoming; Elster, 1979; Schelling, 1984; Thaler and Shefrin, 1981; Winston, 1980).

Magnitude Effects

The effect of magnitude on the discount rate is as strong as the effect of time delay. In both the Thaler and Benzion et al. studies using hypothetical questions, the implicit discount rates declined sharply with the size of the purchase. A similar result has been observed by Holcomb and Nelson (1989) over a small range of actual payoffs, $5–$17. Also, the very high discount rates observed for relatively small hypothetical rewards were obtained by Horowitz (1988) for an actual payoff of $50.

There are two plausible behavioral explanations for the magnitude effect. The first is based on the psychology of perception (psychophysics): people are sensitive not only to relative differences in money amounts, but also to absolute differences (Loewenstein and Prelec, 1989b). The perceptual difference between $100 now and $150 in a year, for example, appears greater than the difference between $10 now and $15 in one year, so that many people are willing to wait for the extra $50 in the first instance, but not for the $5 in the second. The second explanation relies on notions of mental accounting (Shefrin and Thaler, 1988). Suppose that small windfalls are entered into a mental checking account and are largely consumed, while larger amounts are entered into a mental savings account, with a much smaller propensity to consume. Then the cost of waiting for a small windfall may be perceived to be foregone consumption, while in contrast, the opportunity cost of waiting for a large windfall is perceived as simply foregone interest. If foregone consumption is more tempting than foregone interest, the magnitude effect will be observed.[5] (The next installment in this feature, on savings, will discuss these issues in more detail).

Sign Effects

The third strong empirical regularity in the discounting surveys is that the discount rate for gains is much greater than for losses. People are quite anxious to receive a positive reward, especially a small one, but are less anxious to postpone a loss. Part of this preference comes from a simple "debt aversion." Many people pay off mortgages and student loans quicker than they have to, even when the rate they are paying is less than they earn on safe investments.

Reference Points

In descriptive theories of decision making under uncertainty, the distinction between gains and losses has received considerable attention. Decision makers do not appear to integrate outcomes with their wealth or existing consumption level, as normally assumed in expected utility theory. Rather individuals appear to react to events as changes, relative to some natural reference point. This observation was first made by Markowitz (1952), and more recently Kahneman and Tversky (1979) use changes in wealth relative to a reference point as the carriers of value in their prospect theory.

Reference points are also important in intertemporal choice (Loewenstein and Prelec, 1989a). Loewenstein (1988) offers the following demonstration of a reference point effect. An experiment was conducted using 105 high school sophomores and juniors. All subjects received a $7 gift certificate for a local record shop. The expected time at which the students would receive the certificates was varied among one, four,

---

[5] It seems likely that there are also differential discount rates by type of consumption good. One might be more impatient to receive a new car than a new (energy-efficient) furnace, as long as the old furnace works. More research is needed on this question.

Table 1

**Mean Amounts to Speed-up and Delay Consumption
($7 Record Store Gift Certificate)**

Time Interval	Delay	Speed-up	Significance
1 week versus 4 weeks	$1.09	$.25	.001
4 weeks versus 8 weeks	$.84	$.37	.005
1 week versus 8 weeks	$1.76	$.52	.001

and eight weeks. The students were then given a series of binary choices between keeping their certificates at the originally appointed times, or trading them either for smaller certificates to be received earlier, or for larger certificates to be received later. For example, subjects who expected to receive a four week certificate were asked whether they would trade it for an eight week certificate, the value of which was varied between $7.10 and $10.00. They were told that the experimenter would select and implement one of their choices at random.

The design of this experiment allows the role of the reference point to be empirically tested. Some subjects were asked to make a tradeoff between the size of the reward and its *delay* from week 1 to week 4, while other subjects were making a tradeoff between the size of the reward and its *speed-up* from week 4 to week 1. If subjects were not influenced by reference points, then this manipulation would have no effect. The results of the experiment are shown in Table 1. The figures shown are the mean minimum amounts to speed up or delay consumption, depending on the condition. For all three comparisons, the mean delay premium is at least twice the mean speed-up cost, with all differences being statistically significant. Subjects demand more to wait past the expected arrival date than they are willing to pay to speed up its expected arrival. (Similar results are obtained by Benzion et al., 1989.) The result is compatible with Kahneman and Tversky's notion of *loss aversion*, the idea that the disutility of losing a given amount of money is significantly greater in absolute value than the utility of gaining the same amount.

Loss aversion also induces preferences for particular patterns of consumption over time. In situations when past consumption levels set reference points for future consumption, individuals may prefer an increasing consumption profile. For example, Loewenstein and Prelec (1989a) asked 95 Harvard undergraduates three questions. First, the students were asked to choose between two free dinners to be consumed on a Friday night in one month: a dinner at a fancy French restaurant, or a dinner at a local Greek restaurant. Most had the good sense to prefer the French dinner. Then, they were asked whether they would rather have the French dinner in one month or two months. Of those who selected the French dinner originally, 80 percent preferred to have it in one month rather than two, implying a positive discount rate. The third question offered subjects two hypothetical meals, the first in one month, the second in two months. Subjects were asked which order they preferred: Greek in one month, or

French in two months; or French in one month, and Greek in two months. Here, 57 percent of the French food lovers elected to have the Greek meal first. In a standard utility framework, this latter response implies a negative rate of time preference, inconsistent with the answer to the second question. There is no inconsistency, however, if people evaluate current consumption relative to past consumption and are loss averse. They simply prefer a pattern of increasing utility over time.

The preference for a rising consumption profile helps explain an anomaly in labor markets, namely that wages rise with age even when productivity does not (Medoff and Abraham, 1980). In many academic departments, for example, the highest paid faculty are the oldest, even if they are no longer the most productive. The two most important standard explanations for this pattern involve specific human capital and agency costs. The human capital argument is that firms offer the increasing age-earnings profile to encourage workers to stay in the firm long enough to make firm-specific training pay off. The agency cost argument, due to Lazear (1981), suggests that firms offer wages above marginal product for older workers to prevent workers from cheating and shirking. (A worker who gets caught risks losing the present value of the difference between pay and productivity.) While both of these explanations have merit in some occupations (see the articles by Carmichael and by Hutchens in this issue), Frank and Hutchens (forthcoming) show that the same pattern of wages is observed for two occupations in which neither traditional explanation is plausible, namely airline pilots and intercity bus drivers. In the case of pilots, Frank and Hutchens show that wages increase sharply with age while productivity does not. Yet, virtually all the training pilots receive is general, and pilots who shirk on (say) safety are amply punished by nature. Rather, in this case, it seems that the upward sloping age-earnings profile must be due to a preference for income growth, per se.

Evidence for such a pattern of preferences comes from a survey of 100 adults polled at the Museum of Science and Industry in Chicago (Loewenstein and Sicherman, 1989). Respondents were asked to choose between several hypothetical jobs which lasted six years and were identical except in the wage profile they offered. All jobs paid the same total *undiscounted* wages but differed in slope. For one job, wages decreased yearly. For another, they remained constant, and for the remaining five they increased at varying rates. In addition to interest, virtually every economic consideration favored the job with declining wages. For example, if the subject didn't like the job and quit, or was fired before the end of the six years, the declining wage option would provide greater total payments. Despite the incentives for selecting the decreasing wage profile, only 12 percent of the subjects liked it best. Another 12 percent preferred the flat profile, with all other subjects selecting one of the increasing profiles as their favorite.

A result such as this one always makes an economist wonder whether the subjects were just confused. Certainly, if the subjects had the logic of the economic argument explained to them (that the downward sloping wage profile plus saving dominates the others) they would come to their senses, right? To check on this, subjects were asked their preferences again, but after they had been were presented with the economic argument favoring the declining profile, and with psychological arguments in favor of

increasing profiles. The effect of these arguments was minimal. The number of subjects preferring the increasing profile fell from 76 to 69 percent.

The preference for an increasing income stream can be understood by using two concepts discussed above: loss aversion and self-control. Loss aversion explains why workers prefer an increasing *consumption* profile (since the utility of current consumption will depend on previous consumption). Costly self-control explains why workers want an increasing *income* profile, because they cannot rely on themselves to save enough from a flat income (or declining) profile to produce the desired increasing consumption profile.

Savoring and Dread

The standard discounted utility model assumes that the discount rate is constant and, normally, positive. Are there any circumstances in which people prefer to have gains postponed or losses expedited? Marshall (1891, p. 178) suggested one negative influence on the discount rate for gains: "When calculating the rate at which a future benefit is discounted, we must be careful to make allowance for the pleasures of expectation." We will use the terms *savoring* to refer to the positive utility derived from anticipating future pleasant outcomes and *dread* to refer to the negative contemplation of unpleasant outcomes.

The influence of both savoring and dread is demonstrated in the following experiment conducted by Loewenstein (1987). Subjects were asked to specify "the most you would pay now" to obtain (avoid) each of five outcomes, immediately, and following delays of: 3 hours, one day, 3 days, 1 year, and ten years. The five outcomes were: gain $4; lose $4; lose $1000; receive a (non-lethal) 110 volt shock; receive a kiss from the movie star of your choice. The results are plotted in Figure 3.

Discounted utility predicts that the value of a gain and the aversiveness of a loss should decline monotonically with delay before the event occurs. People should want to consume gains as soon as possible and postpone losses as long as possible. As can be seen, however, the two non-monetary outcomes yielded quite different patterns of time preference. For the kiss from the movie star, subjects preferred to delay the outcome for three days, presumably to savor its anticipation. For the electric shock, subjects were willing to pay substantially more to avoid a shock to be received in one or ten years than one in the immediate future. In this case subjects seemed to be willing to pay to avoid having to worry about the event over an extended period of time.

While a kiss from a movie star and an electric shock are rather exotic experiences, Loewenstein (1987) has also obtained similar results for more mundane items. In a demonstration of the utility of savoring, 84 percent of his subjects indicated that they would prefer to receive a dinner at a fancy French restaurant on the second of three weekends rather than the first. To demonstrate dread, subjects were asked: (p. 674) "What is the least amount of money you would accept for cleaning 100 hamster cages at the Psychology Department's animal laboratory. You will be paid immediately.... The job is unpleasant but takes only three hours. How much would you need to be

Figure 3

**Maximum Payment to Obtain / Avoid Outcomes at Selected Times.
Proportion of Current Value (*N* = 30).**
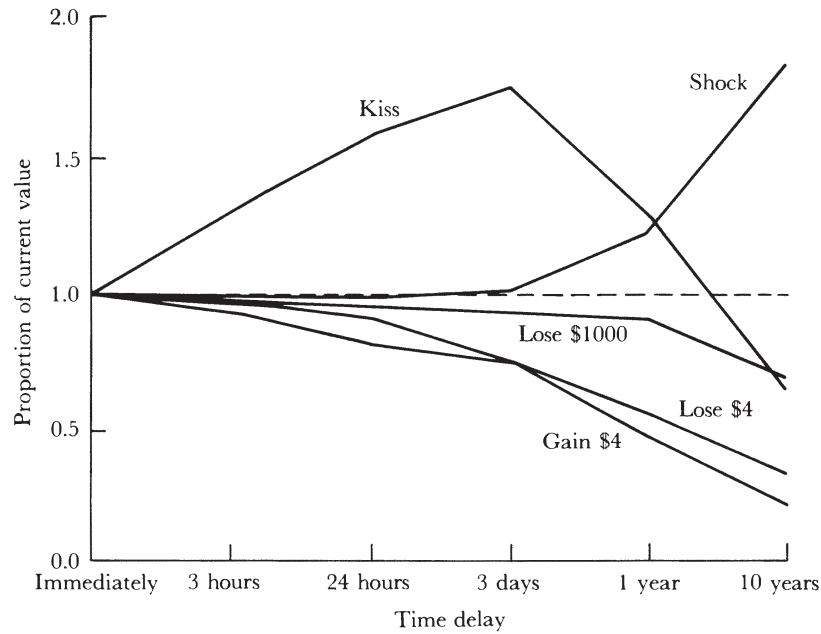
Source: Loewenstein (1987)

paid to clean the cages: (1) once during the next 7 days; (2) once during the week beginning one year from now?" The mean reservation wage for cleaning the cages next week was $30 while the reservation wage for doing the task in a year was $37. In fact, only 2 of 37 subjects gave a smaller response to question (2) than question (1).

Commentary

The policy implications of this line of research are both interesting and treacherous. At a micro level, the high discount rates observed in some contexts (such as appliance purchases) and by some groups (such as teenagers) raise serious questions about consumer rationality. (As mentioned above, in many intertemporal situations involving self-control, individuals question *their own* ability to make rational, long-term choices.) How can it be rational for a consumer to choose a refrigerator that costs $50 less than another equivalent model but consumes $50 more in electricity every year? While such cases do not establish a need for government intervention, the presumption that consumers choose best for themselves is rather weakened.

At a macro level, the psychology of intertemporal choice complicates the already complicated question of selecting the proper social rate of discount. The standard view is that the market rate of interest, corrected for tax distortions, represents an

aggregation of individual time preferences, and is the appropriate social rate of time discounting. However, correcting for tax distortions is far from trivial, and the situation is further complicated by the internationalization of capital markets, which obscures the relationship between time preferences and interest rates in a particular country. Lind (forthcoming) argues that given these complications, the only reasonable way to determine the social rate of time preference is to elicit time preferences at the individual level. But if individuals do not discount everything at a single rate, then which rate is the one that is appropriate for social discounting? Suppose that an individual's freezer purchase implies a discount rate of 50 percent, but that the same person is indifferent between saving 10 lives this year and 10 lives in 20 years? How then should we decide between building another power plant and improving highway safety?

Many economists view the research on the psychology of decision making as a nuisance. The research often provides evidence that individuals violate certain assumptions of rational choice without offering alternative assumptions that can easily be incorporated into economic models. However, psychology can be constructive as well as destructive. For example, in the case of increasing wage profiles, the psychologists' observation that people care about changes in as well as absolute levels of income and consumption (which should be noncontroversial since economists don't argue about tastes) can reconcile the preference for increasing wage profiles with the standard economic assumption that people discount the future. The advantage of drawing on empirical research to suggest modifications in the utility function is that the proposed modifications are less ad hoc. A good example of this kind of reasoning, is offered by Constantinides (1988) in his paper on the "equity premiumpuzzle" (why are returns on stocks so much higher than on bonds?). Constantinides bases his explanation on the assumption that the utility of current consumption depends on past levels of consumption, or as he calls it, habit formation. A cynic might argue that if you try enough utility functions, you can explain anything. However, here that criticism would be misplaced. The habit formation assumption seems to fit intuitions about behavior, and is consistent with a great deal of empirical research. It is even testable. Explanations that rely on assumptions that are testable (or even better, true!) are more attractive than those based on assumptions which are untestable or implausible, for example those which depend on time-varying changes in the unobservable risk of economic catastrophe.

References

Ainslie, George, "Specious Reward: A Behavioral Theory of Impulsiveness and Impulse Control," *Psychological Bulletin*, 1975, *82*, 463–509.

Ainslie, George, *Picoeconomics: The Interaction of Successive Motivational States Within the Individual.* Cambridge: Cambridge University Press, forthcoming.

Benzion, Uri, Amnon Rapoport, and Joseph Yagil, "Discount Rates Inferred from Decisions: An Experimental Study," *Management Science*, March 1989, *35*, 270–84.

Constantinides, George, "Habit Formation: A Resolution of the Equity Premium Puzzle," unpublished working paper, University of Chicago, Graduate School of Business, 1988.

Courant, Paul, Edward Gramlich, and John Laitner, "A Dynamic Micro Estimate of the Life-Cycle Model." In Aaron, Henry G., and Gary Burtless, eds., *Retirement and Economic Behavior.* Washington D.C.: Brookings Institution, 1986.

Elster, Jon, *Ulysses and the Sirens.* New York: Cambridge University Press, 1979.

Frank, Robert, and Robert Hutchens, "Feeling Better vs. Feeling Good: A Life-Cycle Theory of Wages," *Quarterly Journal of Economics*, forthcoming.

Gately, Dermot, "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables: Comment," *Bell Journal of Economics*, 1980, *11*, 373–374.

Hausman, Jerry, "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables," *Bell Journal of Economics*, 1979, *10*, 33–54.

Herrnstein, Richard J., "Relative and Absolute Strength of Response as a Function of Frequency of Reinforcement," *Journal of Experimental Analysis of Behavior.* 1961, *4*, 267–272.

Holcomb, John H., and Paul S. Nelson, "An Experimental Investigation of Individual Time Preference," unpublished working paper, 1989.

Horowitz, John K., "Discounting Money Payoffs: An Experimental Analysis." working paper, Department of Agricultural and Resource Economics, University of Maryland, 1988.

Kahneman, Daniel, and Amos Tversky, "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, 1979, *47*, 363–91.

Lazear, Edward, "Agency, Earnings Profiles, Productivity, and Hours Restrictions," *American Economic Review*, 1981, *61*, 606–620.

Lind, Robert, "Reassessing the Government's Discount Rate Policy in Light of New Theory and Data in a World Economy with a High Degree of Capital Mobility," *Journal of Environmental Economics and Management*, forthcoming.

Loewenstein, George, "Anticipation and the Valuation of Delayed Consumption," *Economic Journal*, 1987, *97*, 666–684.

Loewenstein, George, "Frames of Mind in Intertemporal Choice," *Management Science*, 1988, *34*, 200–214.

Loewenstein, George, and Drazen Prelec, "Anomalies in Intertemporal Choice: Evidence and Interpretation," Russell Sage Foundation Working Paper, 1989a.

Loewenstein, George, and Drazen Prelec, "Decision Making Over Time and Under Uncertainty: A Common Approach," working paper, University of Chicago, Center for Decision Research, 1989b.

Loewenstein, George, and Nachum Sicherman, "Do Workers Prefer Increasing Wage Profiles?" working paper, University of Chicago, Center for Decision Research, 1989.

Marshall, Alfred, *Principles of Economics*, 2nd ed. London: Macmillian, 1891.

Markowitz, Harry, "The Utility of Wealth," *Journal of Political Economy*, 1952, *60*, 151–58.

Medoff, James L., and Katharine G. Abraham, "Experience, Performance and Earnings," *Quarterly Journal of Economics*, December 1980, *95*, 703–36.

Ruderman, Henry, Mark Levine, and James McMahon, "Energy-Efficiency Choice in the Purchase of Residential Appliances." In Kempton, Willett, and Max Neiman, eds., *Energy Efficiency: Perspectives on Individual Behavior.* Washington D.C.: American Council for an Energy Efficient Economy, 1986.

Schelling, Thomas, "Self-Command in Practice, in Policy, and in a Theory of Rational Choice," *American Economic Review*, 1984, *74*, 1–11.

Shefrin, Hersh, and Richard H. Thaler, "The Behavioral Life-Cycle Hypothesis," *Economic Inquiry*, October 1988, *26*, 609–643.

Strotz, Robert H., "Myopia and Inconsistency in Dyanmic Utility Maximization," *Review of Economic Studies*, 1956, *23*, 165–80.

Thaler, Richard H., "Some Empirical Evidence on Dynamic Inconsistency," *Economics Letters*, 1981, *8*, 201–207.

Thaler, Richard H., and Hersh M. Shefrin, "An Economic Theory of Self-Control," *Journal of Political Economy*, *89*, 1981, 392–410.

Winston, Gordon, "Addiction and Backsliding," *Journal of Economic Behavior and Organization*, December 1980, *1*, 295–324.