Title Page:

Encyclopedia of Cognitive Science

Intertemporal Decision Making (article reference code: 708)

David Laibson

Professor of Economics, Department of Economics

Harvard University

Littauer M-14, Department of Economics, Harvard University, Cambridge, Massachusetts

02138

Phone: 617-496-3402; Fax: 617-495-8570

dlaibson@harvard.edu

Headword (title of article): Intertemporal Decision Making

Keywords: discount factor, discount function, discount rate, exponential, hyperbolic, dynamic inconsistency

A-Heads:

- Discounted utility
- Dynamic inconsistency and self-control

Article Definition: Intertemporal decisions imply tradeoffs between current and future rewards. Intertemporal discounting models formalize these tradeoffs, by quantifying the values of delayed payoffs.

1. Introduction:

Decision makers confront a wide range of critical choices that involve tradeoffs between current and future rewards. For example, young workers save part of their paycheck to raise their quality of life in retirement. Habitual heroin users also make decisions with intertemporal consequences when they choose a short-term drug-induced pleasure that jeopardizes their long-term well-being.

To evaluate such tradeoffs, decision makers compare the costs and benefits of activities that occur at different dates in time. The theory of discounted utility provides one framework for evaluating such delayed payoffs. This theory has normative and positive content. It has been proposed as both a description of what people *should* do to maximize their well-being, and to describe what people *actually* do when faced with intertemporal decisions. Both applications of the model are controversial.

Discounted utility models typically assume that delayed rewards are not as desirable as current rewards or similarly that delayed costs are not as undesirable as current costs. This delay effect may reflect many possible factors. For example, delayed rewards are risky because the decision-maker may die before the rewards are experienced. Alternatively, delayed rewards are more abstract than current rewards, and hence a decision-maker may not be able to appreciate or evaluate their full impact in advance. Some contributors have argued that delayed rewards should be valued no less than current rewards, with the sole exception of discounting effects that arise from mortality.

2. Discounted utility

Formal discounting models assume that a consumer's welfare can be represented as a discounted sum of current and future utility. Specifically, the model assumes that at each point in time, t, the decision maker consumes goods $c(t)$. These goods might be summarized by a single consumption index (say a consumption budget for period t), or these goods might be represented by a vector (say apples and oranges). The subjective value to the consumer is given by a utility function $u(c(t))$, which translates the consumption measure, $c(t)$, into a single summary measure of utility at period t.

To evaluate future consumption, the consumer discounts utility with a discount function $F(\tau)$, where τ is the delay between the current period and the future consumption. For example, if the current period is date t and a consumer evaluates consumption half a year from now, the consumer calculates the discounted utility value $F(\frac{1}{2})u(c(t+\frac{1}{2}))$.

Since future consumption is usually assumed to be worth less than current consumption, the discounted utility model posits that $F(\tau)$ is decreasing in τ The more utility is delayed, the less it is worth. Since utility is not undesirable, $F(\tau) \geq 0$ for all values of τ. The model is normalized by assuming $F(0) = 1$. Combining these properties we have

$$1 = F(0) \geq F(\tau) \geq F(\tau') \geq 0,$$

for $0 < \tau < \tau'$. For example if flows of utility a year from now are worth only $\frac{2}{3}$ of what they would be worth if they occurred immediately, then $F(1) = \frac{2}{3}F(0)$.

2.1 Continuous-time discount functions

Intertemporal choice models have been developed in both continuous-time and discrete-time settings. Both approaches are summarized here. Readers without a calculus background may wish to skip directly to the discrete-time analysis.

In continuous-time, the welfare of the consumer at time t — sometimes called the objective function or utility function — is given by

$$\int_{\tau=0}^{\infty} F(\tau)u(c(t+\tau))d\tau,$$

where $F(\tau)$ is the discount function, $u(\cdot)$ is the utility function, and $c(\cdot)$, is consumption.

In both continuous-time and discrete-time models, discount functions are described by two characteristics: discount rates and discount factors. Discount rates and discount factors are normalized with respect to the unit of time, which is usually assumed to be a year.

A discount rate at horizon τ is the rate of decline in the discount function at horizon τ:

$$r(\tau) \equiv \frac{-F'(\tau)}{F(\tau)}.$$

Note that $F'(\tau)$ is the derivative of F with respect to time. Hence, $F'(\tau)$ is the change in F per unit time, so $r(\tau)$ is the rate of decline in F. The higher the discount rate, the more quickly value declines with the delay horizon.

A discount factor at horizon τ is the value of a util discounted with the continuously compounded discount rate at horizon τ:

$$f(\tau) \equiv \lim_{\Delta \to 0} \left(\frac{1}{1+r(\tau)\Delta} \right)^{1/\Delta} = \exp(-r(\tau)).$$

The lower the discount factor, the more quickly value declines with the delay horizon.

2.2 Discrete-time discount functions

Analogous definitions apply to discrete-time models. For this class of models the discount function, $F(\tau)$, need only be defined on a discrete grid of delay values:

$\tau \in \{0, \Delta, 2\Delta, 3\Delta, ...\}$. For example, if the model were designed to reflect weekly observations, then $\Delta = \frac{1}{52}$ years.

Once the discrete-time grid is fixed, the discount function can then be written

$$\{F(0), F(\Delta), F(2\Delta), F(3\Delta), \ldots\}.$$

The welfare of the consumer at time t is given by

$$\sum_{\tau=0}^{\infty} F(\tau\Delta)u(c(t+\tau\Delta)).$$

At horizon τ, the discount function declines at rate

$$r(t) = -\frac{\left(F(t) - F(t-\Delta)\right)/\Delta}{F(t)}.$$

The numerator of this expression represents the change per unit time.

The discount factor at horizon τ is the value of a util discounted with the discount rate at horizon τ compounded at frequency Δ.

$$f(\tau) = \left(\frac{1}{1 + r(\tau)\Delta}\right)^{1/\Delta} = \left(\frac{F(\tau)}{F(\tau-\Delta)}\right)^{1/\Delta}.$$

As the time intervals in the discrete-time formulation become arbitrarily short (i.e. $\Delta \to 0$), the discrete-time discount rate and discount factor definitions converge to the continuous-time definitions.

2.3 Exponential discounting

Almost all discounting applications use the exponential discount function: $F(\tau) = \exp(-\rho\tau)$. This discount function is often written $F(\tau) = \delta^{\tau}$, where $\delta \equiv \exp(-\rho)$.

For the exponential discount function the discount rate is constant and does not depend on the horizon:

$$r(\tau) = \frac{-F'(\tau)}{F(\tau)} = \rho = -\ln \delta.$$

Likewise, the discount factor is also constant:

$$f(\tau) = \exp(-r(\tau)) = \exp(-\rho) = \delta.$$

Figure one plots three discount functions, including an exponential discount function. Note that the exponential discount function displays a constant rate of decline regardless of the length of the delay. Typical calibrations adopt an annual exponential discount rate of 5%.

[Insert Figure 1 about here.]

2.4 Non-exponential discounting

A growing body of experimental evidence suggests that decision-makers' valuations of delayed rewards are inconsistent with the constant discount rate implied by the exponential discount function. Instead, measured discount rates tend to be higher when the delay horizon is short than when the delay horizon is long. One class of functions that satisfy this property are generalized hyperbolas (Chung and Herrnstein 1961; Ainslie 1992; Loewenstein and Prelec 1992). For example,

$$F(\tau) = (1 + \alpha\tau)^{-\gamma/\alpha}.$$

For these functions the rate of decline in the discount function decreases as τ increases:

$$r(\tau) = \frac{-F'(\tau)}{F(\tau)} = \frac{\gamma}{1 + \alpha\tau}.$$

When $\tau = 0$ the discount rate is γ. As τ increases, the discount rate converges to zero. Figure 1 also plots this generalized hyperbolic discount function. Note that the generalized hyperbolic discount function declines at a faster rate in the short-run than in the long-run, matching a key feature of the experimental data.

To capture this qualitative property, Laibson (1997) adopts a discrete-time discount function, $\{1, \beta\delta, \beta\delta^2, \beta\delta^3, \ldots.\}$, which Phelps and Pollak (1968) had previously used to model intergenerational time preferences. This "quasi-hyperbolic function" reflects the sharp short-run drop in valuation measured in the experimental time preference data and has been adopted as a research tool because of its analytical tractability. The quasi-hyperbolic discount function is only "hyperbolic" in the sense that it captures the key qualitative property of the hyperbolic functions: a faster rate of decline in the short-run than in the long-run. This discrete-time discount function has also been called "present-biased" and "quasi-geometric."

The quasi-hyperbolic discount function is typically calibrated with $\beta \simeq \frac{1}{2}$ and $\delta \simeq 1$. Under this calibration the short-run discount rate exceeds the long-run discount rate:

$$r(1) = 1 > 0 = r(\tau > 1).$$

More generally, any calibration with $0 < \beta < 1$, implies that the short-run discount rate exceeds the long-run discount rate.

2.5 Measuring discount functions

Measuring discount rates has proved to be controversial. Most discount rates studies give subjects choices among a wide-range of delayed rewards. The researchers then impute the time preferences of the subjects based on the subjects' laboratory choices.

In a typical discount rate study, the researchers make three very strong assumptions. First, the researchers assume that rewards are consumed when they are received by the subjects. Second, the researchers assume that the utility function is linear in consumption. Third, the researchers assume that the subjects fully trust the researchers' promises to pay delayed rewards.

If these assumptions are satisfied, it is then possible to impute the discount function by offering subjects reward alternatives and asking the subjects to pick their preferred option. For example, if a subject prefers $x today to $y in one year, then the experimenter concludes that

$$F(0) \cdot x > F(1) \cdot y.$$

Once the subject answers a wide-range of questions of this general form, the researcher attempts to estimate the discount function that most closely matches the subject responses.

Such experiments usually reject the exponential discount function in favor of alternative discount functions characterized by discount rates that decline as the horizon is lengthened (Thaler 1981). Related experiments also reject the implicit assumption that the utility function at date t is not affected by consumption at other dates. This "separability" assumption is contradicted by the finding that subjects care both about the level *and* the slope of their consumption profiles (see Loewenstein and Thaler 1989 for findings that violate the discounted utility model).

3. Dynamic consistency and self-control

Exponential discount functions have the convenient property that preferences held at date t do not change with the passage of time. Consider the following illustration of this "dynamic consistency" property. Suppose that at date 0 a consumer with an exponential discount function with discount factor δ, is asked to evaluate a project that requires investments that cost c utils at time 10 with resulting benefits of b utils at time 11. From the perspective of date zero, this project has utility value $\delta^{10}(-c) + \delta^{11}b$. Assume that this utility value is positive and that at date zero the consumer plans to execute the project.

Now imagine that ten periods pass, and the consumer is asked whether she wishes to reconsider her decision to make the planned investment. From the perspective of period ten, the value of the project to the consumer is $-c + \delta b$. The costs are no longer discounted, since they need to be made in the *current* period. Likewise, the benefits are only discounted with factor $F(1) = \delta^1$ since they will now be available only one period in the future.

Note that the consumer's original preference to pursue the project is unchanged by the passage of time, since $\delta^{10}(-c) + \delta^{11}b > 0$ implies that $-c + \delta b > 0$ (divide both sides of the original inequality by δ^{10}). This property of intertemporally consistent preferences is called "dynamic consistency" and the property will always arise when the discount function is exponential. The passage of time will *never* cause the consumer to switch her preference regarding the investment project (unless new information arrives).

However, dynamic consistency is not a general property of intertemporal choice models. In fact, the *only* stationary discount function that generates this property is the exponential discount function. All other discount functions imply that preferences are dynamically inconsistent: preferences will sometimes switch with the passage of time. To see this, reconsider the investment project described above and evaluate it with the quasi-hyperbolic discount function (assume $\beta = \frac{1}{2}$ and $\delta = 1$). Suppose that the project requires an investment that costs 2 utils at time 10 and generates a payoff of 3 utils at time 11. From the time 0 perspective,

$$\beta\delta^{10}(-c) + \beta\delta^{11}b = \frac{1}{2}(-2) + \frac{1}{2}3 = \frac{1}{2},$$

so the project is worth pursuing. However, from the perspective of period 10, the project generates negative discounted utility

$$-c + \beta\delta b = -2 + \frac{1}{2}3 = -\frac{1}{2}.$$

Hence, the project that the consumer wished to pursue from the perspective of time zero, ceases to be appealling once the moment for investment actually arises in period 10.

This example captures a tension that many decision makers experience. From a distance a project seems worth doing, but as the moment for sacrifice approaches the project becomes increasingly unappealing. For this reason, quasi-hyperbolic discount functions have been used to model a wide range of self-regulation problems, including procrastination, credit card spending, and drug addiction.

3.1 Sophistication, commitment and naivite.

The analysis above does not take a stand on whether consumers foresee these preference reversals. Strotz (1956) identifies two paradigms that can be used to analyze the question of consumer foresight: sophistication and naivite.

Sophisticated consumers will anticipate their own propensity to experience preference reversals. Such consumers will recognize the conflict between their early preference — i.e., the preference to undertake the investment project — and their later contradictory preferences. Such sophisticated consumers may look for ways to lock themselves into the investment activity. For example, consider a person who forces himself to exercise by making an appointment with an expensive trainer.

At the other extreme, Strotz also considered consumers who exhibit naivite about their future preference reversals. Such consumers fail to foresee these reservals and expect to engage in investments that they will not actually carry out (e.g., quitting smoking or completing a project with no deadline). Akerlof (1991) discusses such procrastination problems and O'Donoghue and Rabin (2001) propose a framework that continuously indexes the degree of naivite.

4 Summary

The discounted utility model provides a way of formally evaluating intertemporal tradeoffs. The principle component of the model is a discount function that is used to calculate the discounted value of future utility flows. The key characteristics of the discount function are the discount rate and the discount factor. The discount rate measures the rate of decline of the discount function. The discount factor measures the value of a discounted util. Exponential discount functions are commonly used in most applications of the discounted utility model. Exponential discount functions have a constant discount rate. Exponential discount functions also have the convenient property that they do not generate preference reversals. However, the experimental evidence contradicts the constant discount rate property. Most experimental evidence suggests that the discount rate declines with the length of the delay horizon. Such discounting patterns may play a role in generating self-control problems.

References

Ainslie G (1992) *Picoeconomics.* Cambridge: Cambridge University Press.

Akerlof GA (1991) Procrastination and Obedience. *American Economic Review Papers and Proceedings.* 81, 1-19.

Chung SH and Herrnstein RJ (1961) Relative and Absolute Strengths of Response as a Function of Frequency of Reinforcement. *Journal of the Experimental Analysis of Animal Behavior.* 4, 267-72.

Laibson DI (1997) Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics.* 62, 443-478.

Loewenstein G and Prelec D (1992) Anomalies in Intertemporal Choice: Evidence and an Interpretation. *Quarterly Journal of Economics.* 57, 573-98.

Loewenstein G and Thaler RH (1989) Anomalies: Intertemporal Choice, *Journal of Economic Perspectives.* 3, 181-93.

O'Donoghue T, and Rabin M (2001) Choice and Procrastination. *Quarterly Journal of Economics.* 66, 121-160.

Phelps ES and Pollak RA (1968) On Second-best National Saving and Game-equilibrium Growth. *Review of Economic Studies.* 35, 185-199.

Strotz RH (1956) Myopia and Inconsistency in Dynamic Utility Maximization. *Review of Economic Studies.* 23, 165-180.

Thaler RH (1981) Some Empirical Evidence on Dynamic Inconsistency. *Economic Letters.* 8, 201-207.

Bibliography

Angeletos G, Laibson DI, Repetto A, Tobacman J, Weinberg S (2001) The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation. *Journal of Economic Perspectives*, 15(3), 47-68.

Frederick S., Loewenstein G and O'Donoghue T (forthcoming). Time Discounting and Time Preference: A critical review.. *Journal of Economic Literature.*

Kirby KN (1997) Bidding on the Future: Evidence Against Normative Discounting of Delayed Rewards. *Journal of Experimental Psychology.* 126, 54-70.

Laibson DI (2001) A Cue-Theory of Consumption, *Quarterly Journal of Economics.* 66, 81-120.

Loewenstein G (1996) Out of control: visceral influences on behavior. *Organizational Behavior and Human Decision Processes*, 65, 272-92.

Loewenstein G and Prelec D (1991) Negative time preference, *American Economic Review Papers and Proceedings,* 82, 347-352.

Loewenstein G, Read D and Kalyanaraman S (1999) Mixing virtue and vice: The combined effects of hyperbolic discounting and diversification. *Journal of Behavioral Decision Making*, 12, 257-73.

O'Donoghue T and Rabin M (1999). Doing It Now or Later. *American Economic Review.* 89, 103-124.

Loewenstein G and Elster J editors (1992) *Choice Over Time.* New York: Russell Sage Foundation Press.

Loewenstein G, Read D and Baumeister R editors (forthcoming) *Intertemporal Choice.* New York: Russell Sage Foundation Press.

Thaler, RH and Shefrin HM (1981) An Economic Theory of Self-Control. *Journal of Political Economy.* 89, 392-410.

Glossary:

discount factor: the value of a util discounted with a continuously compounded discount rate.

$$\text{discount factor} \equiv \lim_{\Delta \to 0} \left(\frac{1}{1 + r\Delta} \right)^{1/\Delta}.$$

discount function: weighting function used to calculate the discounted value of future utility flows.

discount rate: rate of decline in the discount function, $F(\cdot)$.

$$\text{discount rate} \equiv \frac{-F'}{F}.$$

dynamically inconsistent time preferences: discount function that generates preference reversals with the passage of time.

exponential discount function: $F(\tau) = -\exp(\rho\tau) = \delta^{\tau}$.

hyperbolic discount function: $F(\tau) = (1 + \alpha\tau)^{-\beta/\alpha}$.

quasi-hyperbolic discount function: $\left\{1, \beta\delta, \beta\delta^2, \beta\delta^3, ...\right\}$.

Illustrations: see attached figure 1.

Word processing package: Scientific Workplace (Latex)

Figure 1: Three Discount Functions