Journal of Conflict Resolution
http://jcr.sagepub.com

Time Horizons, Discounting, and Intertemporal Choice

Philip Streich and Jack S. Levy

Journal of Conflict Resolution 2007; 51; 199
DOI: 10.1177/0022002706298133

Journal of Conflict Resolution
Volume 51 Number 2
April 2007 199-226
© 2007 Sage Publications
10.1177/0022002706298133
http://jcr.sagepub.com
hosted at
http://online.sagepub.com

Time Horizons, Discounting, and Intertemporal Choice

Philip Streich
Jack S. Levy
Department of Political Science
Rutgers University, New Brunswick, New Jersey

Although many decisions involve a stream of payoffs over time, political scientists have given little attention to how actors make the required tradeoffs between present and future payoffs, other than applying the standard exponential discounting model from economics. After summarizing the basic discounting model, we identify some of its leading behavioral anomalies—declining discount rates; preference reversals; higher discount rates for smaller payoffs than for larger payoffs and for gains than for losses; framing effects based on expectations; and a preference for ascending rather than descending sequences. We examine the leading alternative models of discounting and then apply a quasi-hyperbolic discount model to the problem of cooperation in iterated Prisoner's Dilemma games. We demonstrate that if actors display the widely observed tendency to highly discount the immediate future, then cooperation in an iterated Prisoner's Dilemma game is more difficult than Axelrod suggests.

Keywords: *discounted utility; exponential discounting, hyperbolic discounting; quasi-hyperbolic discounting; intertemporal choice; iterated Prisoner's Dilemma*

Many of the topics of interest to students of politics, from economic cooperation to environmental policy to preventive war, involve a stream of payoffs to actors over time. This raises the questions of how much actors value the future relative to the present and what kinds of tradeoffs actors make between current payoffs and future payoffs. Decisions regarding preventive war against a rising adversary, for example, are based on tradeoffs between the costs of a war a state expects to win now and the costs and risks associated with continued decline, a loss of bargaining leverage, and the risk of war under less favorable circumstances later. The greater the weight a declining actor gives to the future, the greater its incentives for war now (Levy 2007). To take another example, the ability to sustain cooperation in an iterated Prisoner's Dilemma game depends on actors having a sufficiently large discount factor or "shadow of the future" (Axelrod 1984, 15).

Authors' Note: We have benefited from the comments and suggestions of Jeffrey Ritter and audiences at our panels at meetings of the Peace Science Society (2004), the Midwest Political Science Association (2006), and the New York Area Political Psychology Group (2006).

Although political scientists recognize in principle that people discount the future and make tradeoffs over time, they give little if any attention to the question of how people discount and how they make tradeoffs between the present and the future. By contrast, economists have devoted considerable attention to the question of "intertemporal choice" involving tradeoffs among payoffs occurring at different points in time. Analyses go back to the early nineteenth century (Frederick, Loewenstein, and O'Donoghue 2003, 15-18), but the study of intertemporal choice changed significantly after the publication of Paul Samuelson's (1937) "discounted utility" (DU) model. This DU model is now standard in economics and among political scientists applying economic models.

One key assumption of the DU model is that discounting is constant over time—people have a single discount factor they apply to each period of time. If a good is worth 10 percent less to you a year from now than it is today, it will be worth another 10 percent less the following year, another 10 percent less the following year, and so on. The assumption of constant-rate discounting has the advantages of generating an exponential discounting function that reflects the opportunity cost of money (which compounds at an exponential rate) and leading to consistent behavior over time, and it has the nice mathematical property of converging over time.

Whether the DU model and its key assumption of a constant-rate discount factor are descriptively accurate is a question that has generated greater debate among both experimental economists and social psychologists. Experimental research has uncovered a series of behavioral anomalies in the DU model, and scholars have constructed a number of alternative models of intertemporal choice that might account for some of these anomalies. Our aim is to review this literature and then apply its key findings to an important problem in cooperation theory.

We begin by summarizing the DU model and its key underlying assumptions. We explore the leading anomalies in the DU model, as revealed by experimental and field studies, and then examine a set of alternative "hyperbolic" and "quasi-hyperbolic" discount models that have been designed to account for some of these anomalies. We end with an analysis of Axelrod's (1984) model of cooperation in iterated Prisoner's Dilemma games, a model that incorporates the DU model for iterated payoffs. We demonstrate that if players behave as quasi-hyperbolic discounters rather than exponential discounters, as observed in experimental studies of intertemporal choice on other issues, then cooperation in iterated Prisoner's Dilemma games is more difficult than Axelrod's model implies.

The Discounted Utility Model

Intertemporal preferences, or preferences over outcomes that occur at different points in time, are typically represented by the DU model. The DU model multiplies an individual's utility in each period by a weight that represents the individual's rate of time preference and aggregates the sum of the weighted utilities:

$$U(x) = \sum_{t=0}^{T} D(t)u(x_t),$$

where $U(x)$ is the discounted present value utility function, $u(x_t)$ is the cardinal instantaneous utility function for payoff x, and $D(t)$ is the discount function. $U(x)$ is the sum in the present period of all of the discounted present values of future utilities that will result from a decision today that generates payoffs T periods into the future. Nearly all discount functions are monotonically decreasing, reflecting the idea that future outcomes are generally worth less today than at the time they occur in the future.

The discount function can take any form, but in the DU model the assumption of constant-rate discounting from one period to the next generates an exponential discount function, $D(t) = \delta^t$, where δ is the discount factor. If we simplify by using payoffs x_i in place of utility, assuming that utility is linear in payoffs, we get

$$U(x) = x_0 + \delta x_1 + \delta^2 x_2 + \ldots + \delta^T x_T = \sum_{t=0}^{T} \delta^t x_t,$$

where $\delta^i x_i$ is the discounted present value of payoff x_i from period i. The exponential discount function for $\delta = .9$ is portrayed in Figure 1.

The discount factor in exponential, constant-rate discounting is generally assumed to range from zero to one. The former represents a future that is entirely discounted and has no present value, and the latter represents a situation in which there is no discounting, and the future is valued as much as the present. A discount factor greater than one would mean that the future is valued more than the present and imply that individuals delay consumption.

Economist Paul Samuelson (1937) first proposed the DU model seven decades ago in an article (his first) titled "A Note on Measurement of Utility." Samuelson wanted to show that the use of consumption over time could facilitate the measurement of the utility of monetary income. The DU model makes six assumptions:

1. *Positive time preference*. People prefer to receive goods sooner rather than later. Thus, the discount factor is less than one.
2. *Constant-rate discounting*, also referred to as time consistency or time invariance. The ratio of the discount functions of any period and the preceding period is simply the discount factor, $\delta^{t+1}/\delta^t = \delta$. Any preference relation between separately occurring outcomes is unaffected if all outcomes are delayed or accelerated by an equal length of time. If an individual prefers one apple today to two apples tomorrow, then she will prefer one apple in one year to two apples in a year and a day.

Figure 1
Exponential Discount Function $D(t) = \delta^t$, with $\delta = .9$

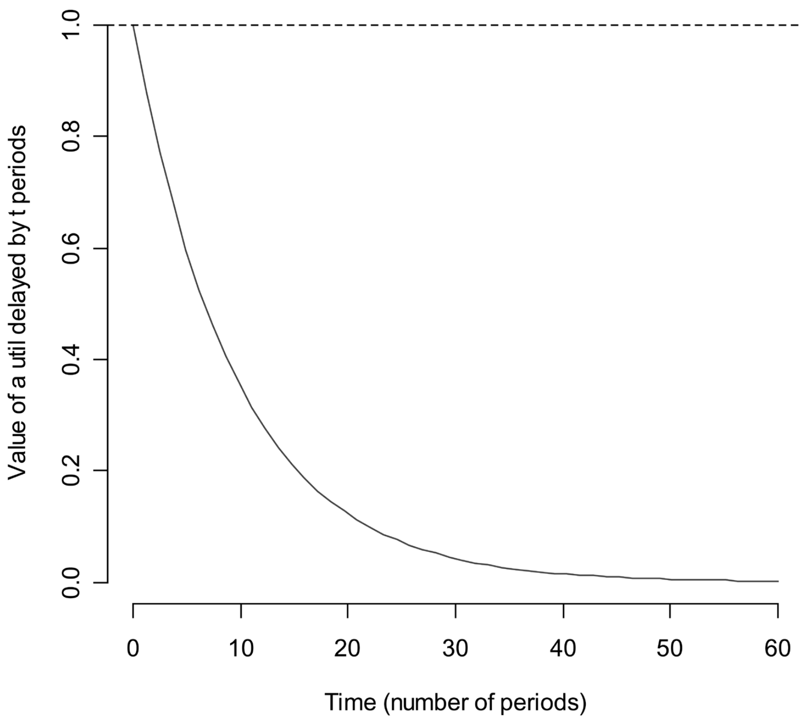

3. *Independence of utility*. Utility in each period is independent of utility in every other period. Utilities are separable, so that the intertemporal utility function, $U(x)$, can be calculated from the sum of every future period's discounted utility (Samuelson 1937, 156; Frederick, Loewenstein, and O'Donoghue 2003, 21).
4. *Independence of consumption*. Preference orderings over consumption or outcomes are not affected by the outcome of any one period. Having an ice cream two days in a row should not affect my preference for ice cream over other desserts the next day.
5. *Stationary instantaneous utility*. The utility function is constant over time. Individuals do not change their preferences as they move through the sequence of choices.
6. *Independence of discounting*. This assumes that the discount function does not vary over types of consumption. Thus, discounting is the same for money as it is for apples or for anything else (Frederick, Loewenstein, and O'Donoghue 2003, 21-22).[1]

A key characteristic of the DU model is its mathematical tractability. Under exponential discounting, the sum of the terms of an infinite series (such as the one that results from an infinitely iterated Prisoner's Dilemma game) will converge. This means that the analyst can calculate the sum of an infinite series of repeated outcomes:

$$\sum_{t=0}^{\infty} \delta^t = 1 + \delta + \delta^2 + \delta^3 + \ldots = \frac{1}{1-\delta}.$$

It is important to note that many behavioral economists and social psychologists often refer to the "discount rate" rather than to the "discount factor" used by most economists and political scientists. The discount factor is a weight and represents the proportion of the utility in period *i* that is incorporated into the discounted present value. The higher the discount factor, the more the future is valued.[2] The discount rate reflects the commonsense notion that the more the future is "discounted," the less it is valued and the lower its discounted present value.

The discount rate k can be calculated from the discount factor δ as follows:

$$\delta = \left(\frac{1}{1+k}\right), \text{ or } k = \frac{1}{\delta} - 1.$$

Thus, the exponential discount function, $D(t)$, in the DU model is related to the discount rate as follows:

$$D(t) = \delta^t = \left(\frac{1}{1+k}\right)^t.$$

Behavioral Anomalies in the DU Model

In proposing his influential DU model, Samuelson (1937) recognized that the model might not fully and accurately represent the way people discount the future and make intertemporal choices. On more than one occasion, he cited the "arbitrariness" of his assumptions and the "serious limitations" of his exercise, which "almost certainly vitiate it even from a theoretical point of view" (pp. 156, 159). He concluded the article by stating that "any connection between utility as discussed here and any welfare concept is disavowed" (p. 161). In subsequent work, Samuelson (1939; 1952, 674) reinforced his concerns about the DU model's normative and descriptive validity: "This [the 1937 article] was intended only as an axiomatic experiment, an intellectual curiosity, which served to reveal to my own satisfaction the arbitrariness of the assumptions and the barrenness of the results" (Samuelson 1939, 291-2).

Although Samuelson (1937, 156) viewed his discounting model as based on assumptions that were analytically useful but not necessarily descriptively accurate, economists quickly accepted the DU model into the formal modeling mainstream of microeconomics and then game theory. As O'Donoghue and Rabin (1999) note, the economics profession has "evolved from perceiving [the DU model] as a useful, ad hoc approximation of intertemporal choice behavior," to errantly "perceiving it as a fundamental axiom of (rational) human behavior" (p. 103).

In the last few decades, experimental research on intertemporal preferences in social psychology and behavioral economics has identified a number of limitations of the generalized DU model. Attempts to measure discount factors for money and certain goods found that discount factors vary greatly with the type of good and the length of delay of the payoff and also with age (Ainslie and Haendel 1983; Loewenstein 1987; Kirby and Marakovic 1995; Kirby and Marakovic 1996, 102; Chapman 1996; Ganiats et al. 2000; Frederick, Loewenstein, and O'Donoghue 2003, 47-50). In addition, some of these studies have found discount factors greater than one, reflecting a negative time preference in which a particular good or outcome is valued more highly in the future than in the present. Studies have also found discount factors that are effectively zero, suggesting that individuals attribute little or no value to future goods (Frederick, Loewenstein, and O'Donoghue 2003, 47-50).[3] For most individuals and most goods, however, the discount factor falls into the (0,1) range.

Of greater interest for our purposes are six anomalies that directly violate some of the axiomatic assumptions of the discount utility model, commonly labeled "declining discount rates," the "common difference effect," the "absolute magnitude effect," "gain-loss asymmetry," "delay-speedup asymmetry," and "preferences for improving sequences."[4] The first two anomalies bear on the DU model's assumption of constant-rate discounting. The third and fourth anomalies relate to the assumption that discounting is independent of the utility of the payoffs. The fourth and fifth anomalies represent framing effects, and suggest that different verbal representations of mathematically identical problems can generate different preferences. The sixth anomaly contradicts the assumption of positive time preference.

Declining Discount Rates

Experimental evidence suggests that discount rates tend to decline as one looks further into the future. The greater the delay between the present (the time of decision) and the point at which the payoff occurs, the lower the discount rate, and hence, the greater the discount factor and the relative weight attached to the outcome (Thaler 1981; Benzion, Rapoport, and Yagil 1989, 277-278; Pender 1996, 282; Frederick, Loewenstein, and O'Donoghue 2003, 24-25; Frederick 2003). This is the "the most robust conclusion" from the literature on intertemporal choice (O'Donoghue and Rabin 2003, 218). What this means is that the discount function flattens out more than the exponential DU model predicts. People are more patient for more temporally distant rewards, relative to those who engage in constant-rate discounting, and they will accept a disproportionately lower amount of compensation for longer delays for forgoing a reward in the present.

In one experimental study, subjects were asked to name the amount of money that would leave them indifferent to receiving $100 in the present and a comparable amount in one month, one year, and ten years. The median responses were \$150 for one year, \$500 for five years, and \$1,000 for ten years, representing discount rates of 50 percent, 38 percent, and 26 percent, respectively (Frederick 2003, 100).[5] Other studies generate comparable results. A particularly important aspect of declining discount rates is the finding observed in most experimental and field studies that discount rates on payoffs in the very near future—especially the period after the present—are quite high. That is, the value of a reward drops significantly in the immediate future, so that individuals are very impatient with regard to short time delays (Laibson 1997, 445; Angeletos et al. 2003, 519).

Common Difference Effect or "Dynamic Inconsistency"

As a consequence of declining discount rates, an individual's preference between two outcomes, one occurring a given time after the first, may not be consistent when each is delayed by an equal length of time (common difference). Contrary to the assumption of constant-rate discounting, an individual may prefer x now to a discounted y in the next period, but when x and y are each delayed t periods and each is discounted, he may prefer y in $t+1$ periods from now over x in t periods. For instance, an individual might prefer one apple today to two apples tomorrow, yet prefer two apples 366 days from now to one apple 365 days from now. This dynamic inconsistency contradicts the exponential discounting model's assumptions of time consistency and constant-rate discounting. The resulting preference reversals are inconsistent with standard rational models of behavior and problematic for any normative theory of choice.[6]

In one example, respondents in a survey are given a choice between \$100 today or \$115 next week. Most of the respondents chose the \$100. Then, the respondents were asked whether they preferred \$100 in fifty-two weeks or \$115 in fifty-three weeks. Most of the respondents chose the latter (Herrnstein 1990, 358). Thus, preferences for \$100 over \$115 a week later were reversed when the common delay length of one year was added to each reward. Studies with different rewards produce similar results (Kirby and Herrnstein 1995, 83). A constant-rate exponential discount function would maintain preference for the smaller, earlier reward over the larger, later reward after a delay of a year was added to each reward. Note that the example reflects a discount rate that is high in the near future but lower in the more distant future, which follows the general finding of declining discount rates over time.

Relatively high short-term discount rates, in conjunction with subsequently declining discount rates, explain these preference reversals.[7] If the present value of an apple tomorrow is only 40 percent of the value of an apple today, then one would prefer an apple today to two apples tomorrow. The value of apples tomorrow but not the one today is discounted. Many periods into the future, however, the values of the apples on successive days are both discounted, the second only slightly more than the latter, given declining discount rates. Hence, today one would prefer two apples a year and a day from now to one apple a year from now.

Absolute Magnitude Effect

Individual discount rates are lower for large payoffs than they are for small payoffs so that as the value of rewards increases, the discount rate of individuals generally decreases (Ainslie and Haendel 1983; Loewenstein 1987; Benzion, Rapoport, and Yagil 1989, 278; Kirby and Marakovic 1995; Kirby 1997). This means that people give greater weight to large future payoffs than to smaller future payoffs. For example, in one study, respondents were on average indifferent between $15 now and $60 in one year, between $250 now and $350 in one year, and between $3,000 now and $4,000 in one year (Thaler 1981, 204). The discount rates for these pairs are 300 percent, 40 percent, and 33 percent, respectively.[8] This violates the DU model's assumption of independence of discounting from consumption, which requires that the discount rate be independent of the value of the outcome.

Gain-Loss Asymmetry

People tend to discount future gains more than they do future losses so that in relative terms, people give more weight to future losses than to future gains of the same magnitude. For example, respondents in one study were indifferent between receiving $10 now and $21 in one year and indifferent between losing $10 now and losing $15 in one year (Loewenstein and Prelec 1992, 575). The first indifference pair involves gains and reflects a discount rate of 110 percent, while the second indifference pair involves losses and reflects a discount rate of 50 percent.[9] Other studies produce similar findings (Loewenstein 1987; Benzion, Rapoport, and Yagil 1989, 278). In this anomaly, as in the absolute magnitude effect, the type of outcome is affecting the discounting. Both violate the independence of discounting from consumption assumption, in which it is assumed that the type of good or outcome being discounted does not influence the discount rate.

Delay-Speedup Asymmetry

In the change in value people attach to outcomes, it makes a difference whether they receive a given outcome sooner than they expect or later than they expect. Discount rates can vary widely depending on whether a choice problem is portrayed as a delay or acceleration in the receipt of a payoff. One study hypothetically offered a VCR worth $300 to respondents but varied expectations as to when respondents would receive the reward: some thought they would receive the reward immediately, while others thought they would receive the reward in a year. Respondents expecting to receive the reward in a year were then asked how much they would pay to receive the VCR immediately, while respondents expecting to receive the VCR immediately were asked how much they wanted in return for delaying the receipt of the VCR for a year. On average, respondents expecting to receive the VCR in a year were willing to pay $54 to receive it immediately, but those who thought they were receiving the VCR immediately demanded over $126 to allow the payoff to be delayed for a year (Loewenstein 1988, 207-8). Benzion, Rapoport, and Yagil (1989, 278) get similar results. Thus, the framing of an outcome relative to a reference point is consequential, contrary to the rationalist assumption built into the DU model, that mathematically identical choice problems should generate identical choices.

The gain-loss effect and delay-speedup effect are reminiscent of anomalies in expected utility theory that scholars have integrated into prospect theory (Kahneman and Tversky 1979). Each of these anomalies derives from reference dependence, framing effects, and the asymmetrical impact of losses and gains.[10] The greater weight given to future losses than to future gains is consistent with the phenomenon of loss aversion in prospect theory—the idea that losses bring more pain than gains bring pleasure and that losses are overweighted relative to comparable gains. Similarly, in the delay-speedup anomaly, if receiving a reward sooner than expected is treated as a gain and if receiving a reward later than expected is treated as a loss, it is not surprising, in light of prospect theory and the overweighting of losses, that people will demand more to be compensated for a loss than they will pay to secure a comparable gain.

Preference for Improving Sequences

When payoffs are presented as sequences of descending or ascending value, there is a strong pattern of preference for improving sequences, contrary to the DU model and to any model of positive time preference (Loewenstein 1987, 1992; Elster and Loewenstein 1992; Loewenstein and Sicherman 1991; Loewenstein and Prelec 1993). There is some evidence, for example, that individuals favor improving sequences when considering salaries and contract payments, even when the discounted presented value of that sequence is less than that for a series of decreasing outcomes (Loewenstein and Sicherman 1991; Hsee, Abelson, and Salovey 1991).

In one widely cited experiment, Loewenstein and Prelec (1993, 93) asked respondents to choose between sequences of free dinners at French and Greek restaurants with a time delay between the meals. First, 86 percent of the respondents stated that they preferred the French restaurant to the Greek in a direct comparison. When those who preferred the French restaurant were asked whether they preferred a free dinner at a Greek restaurant next week together with a free dinner at a French restaurant in one month or the dinner at the French restaurant next week and the dinner at the Greek restaurant in one month, the majority of respondents chose to eat at the Greek restaurant first and the French restaurant later.

The preference for improving sequences contradicts the fundamental assumption of the DU model that people have positive time preferences and that they prefer positive payoffs sooner rather than later. It also violates the assumption of independence of utility, since a preference over a sequence of utilities implies that the individual is not separating the utilities but considering them in relation to each other (Frederick, Loewenstein, and O'Donoghue 2003, 29).

The negative time preference associated with improving sequences is often attributed to savoring or anticipation effects. The process of waiting and thinking about a positive future payoff generates an anticipation effect that itself creates a positive utility (Loewenstein 1987). If future payoffs are negative, people experience a negative utility from the anticipation or dread of the expected event. That is, people experience utility both directly from an event itself and also from the knowledge of the impending event (Loewenstein 2006).

Further evidence of anticipation effects emerges from recent research in neuroscience, which permits more direct measurement of individuals' thoughts and feelings.[11] In one experiment, Berns et al. (2006) found that subjects who knew they would be getting electric shocks preferred receiving those shocks sooner rather than later. They also preferred stronger shocks sooner to lesser shocks later.

An alternative interpretation of preference for improving sequences is based on reference dependence. If the first payoff in a sequence serves as a reference point, then subsequent payoffs in a declining sequence are defined as losses and overweighted, while future payoffs in an improving sequence are defined as gains.[12]

The accumulation of experimental research on intertemporal choice has made it increasingly clear that the exponential discounting model that Samuelson (1937) pioneered nearly seventy years ago, which has subsequently dominated economics and economic applications in political science, does not provide a descriptively accurate model of how most people actually behave in making choices over time. Instead of discounting by a constant rate from one period to the next, people tend to discount relatively more heavily the near-term future and to discount relatively less heavily the more distant future, compared to constant-rate exponential discounting. In addition, discounting is not independent of the value of future outcomes. People have greater discount rates for less valuable outcomes than they do more valuable outcomes, and they have greater discount rates for gains than for losses. This asymmetry of losses and gains, so familiar to students of prospect theory, carries over into other manifestations of reference dependence and framing: the anticipated loss of utility of having to wait longer than expected for a future reward is greater than the anticipated gain in utility from receiving a future reward sooner than expected.

Observational evidence of anomalies in the exponential discounting function has led researchers to search for alternate discounting functions that provide better descriptions of intertemporal choice behavior. They have discovered no magic formula that accounts for all anomalies in the DU model, but they have suggested several alternative models that are worth exploring further. The most basic is hyperbolic discounting, which incorporates the tendency for individuals to heavily discount the immediate future while giving more weight to the distant future, relative to the standard DU model.

Hyperbolic Discounting Models

Hyperbolic models are appealing for the analysis of discounting because they incorporate two key patterns of behavior observed in experimental studies: relatively steep discounting in the near present and a declining rate of time preference, which generates a flattening out of the discount function over time relative to the exponential model. Economist Robert Strotz (1956) is probably the first to formalize a theory of commitment mechanisms that utilized a nonexponential discount function.[13] Experimental research with animal subjects (mostly pigeons) in the 1960s and 1970s (Chung and Herrnstein 1961; Rachlin and Green 1972; Ainslie 1974; Mazur 1987) was also instrumental in developing hyperbolic models.[14] Although the DU model still dominates economics, hyperbolic discounting is increasingly accepted. As Prelec (2004) says, "Few economic hypotheses have advanced so rapidly from the fringe to the mainstream as hyperbolic discounting" (p. 511).

The most common hyperbolic discount functions proposed by researchers are the following (Frederick, Loewenstein, and O'Donoghue 2003, 69):

1. $D(t) = 1/t$, where t equals the length of delay (Ainslie 1975).
2. $D(t) = 1/(1 + kt)$, where k equals the discount rate (Herrnstein 1981; Mazur 1987).
3. $D(t) = 1/(1 + \alpha t)^{\gamma/\alpha}$, the generalized hyperbolic discount function, with α, $\gamma > 0$, where α determines "how much the function departs from constant discounting" (Loewenstein and Prelec 1992, 580) and γ is a time preference parameter positively related to the instantaneous discount rate (Laibson et al. 1998, 100).[15]

Each of the three hyperbolic discounting functions is designed to account for the anomalies of declining discount rates and common difference effects. They cannot account for anomalies relating to preferences for improving sequences, to effects relating to the magnitude of the payoff, or to framing effects that derive either from asymmetries between losses and gains or from the differences between the delay or acceleration in the receipt of a good.[16]

Figure 2 shows an exponential discount function, $D(t) = \delta^t$, with $\delta = .9$, and the generalized hyperbolic discount function suggested by Loewenstein and Prelec (1992), $D(t) = (1 + \alpha t)^{-\gamma/\alpha}$, with $\alpha = 2$ and $\gamma = 1$.

While the third discount function, proposed by Loewenstein and Prelec (1992), probably provides the best fit with the experimental data, it is quite complex, and many students of intertemporal choice prefer the second function because it is more parsimonious (Rachlin and Raineri 1992, 98; Green, Fry, and Myerson 1994, 33; Henderson and Langford 1998, 1494).

Figure 2
Exponential and Hyperbolic Discount Functions
$D(t) = \delta^t$, $\delta = .9$; $D(t) = (1 + \alpha t)^{-\gamma/\alpha}$, $\alpha = 2$, $\gamma = 1$

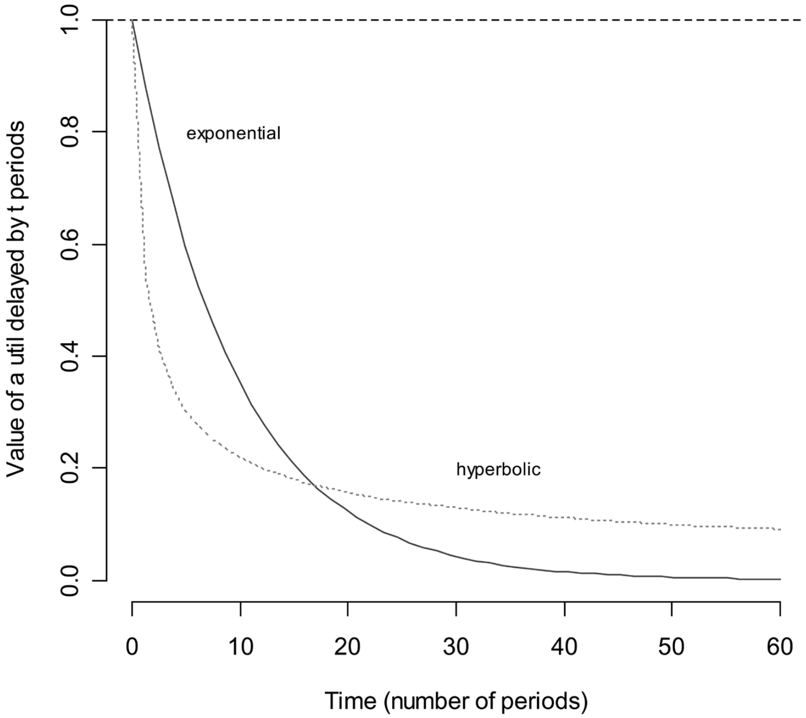

This and other hyperbolic functions are most commonly used to calculate the discounted present value of a single outcome in the future or perhaps a handful of such outcomes, not of an infinitely iterated series. Loewenstein and Prelec (1992), for example, state that their hyperbolic discounting model with a modified value function is “primarily concerned with explaining elementary types of intertemporal choice involving no more than two or three distinct dated outcomes” (p. 578). The primary reason for this is the fact that many hyperbolic discount functions do not converge. In an infinite series of payoffs, such as those in an iterated prisoner’s dilemma, the sum of a hyperbolic discount function continues to grow unbounded as t goes to infinity. The divergent property of the hyperbolic discount function is a serious limitation, particularly given the tractability of the convergent exponential DU model.[17] This problem has led researchers in economics to explore the discrete-time “quasi-hyperbolic discount function” (Laibson 1997; Laibson et al. 1998) in an attempt to fit the data with a function that is both descriptively accurate and mathematically tractable.

The Quasi-Hyperbolic Discounting Model

First used by Phelps and Pollak (1968, 186), and later by Laibson (1997; Laibson et al. 1998; see Barro 1999), the quasi-hyperbolic discount function is expressed as the set of discrete values $\{1, \beta\delta, \beta\delta^2, \beta\delta^3, \ldots\}, 0 < \beta \leq 1$. The first term (1) is the weight for the present period, $\beta\delta$ the weight for the first subsequent period, and so on. When expanded, this model is quite similar to the DU model:

$$x_0 + \beta\delta x_1 + \beta\delta^2 x_2 + \ldots + \beta\delta^T x_T = x_0 + \sum_{t=1}^{T} \beta\delta^t x_t.$$

If β equals 1, then the quasi-hyperbolic model simply reduces to the DU model. The change in value of a payoff from the present to the first period in the future is $\beta\delta/1 = \beta\delta$. In any period after the first, however, the successive per period change in value of a payoff is $\beta\delta^{t+1}/\beta\delta^t = \delta$. Thus, the quasi-hyperbolic discount function is similar to exponential discounting after the first period. When $\beta < 1$, as is generally observed (Laibson 1997, 452), then $\beta\delta < \delta$. The discounted present value of the payoff in the next period after the present is less in the quasi-hyperbolic model than in the standard DU model.

The parameter β refers to the high discounting individuals apply to payoffs in the next period after the present (giving it lesser weight) and is sometimes referred to as the need for "immediate gratification" (O'Donoghue and Rabin 2003, 230). After the period after the present, quasi-hyperbolic discounters engage in exponential discounting. Researchers have generally found that the quasi-hyperbolic discount function roughly approximates discounting behavior observed in the laboratory, particularly in the first few periods. Researchers have empirically estimated β to fall in the .60 to .70 range but, like the discount rate, this parameter can vary widely (Angeletos et al. 2003, 522; Laibson 1997, 452).

The quasi-hyperbolic discounting model is only partially hyperbolic. It captures the steep drop in the discounted present value when one moves from the present to the next period and hence the high impatience for tradeoffs with the immediate future. In doing so, it can also account for preference reversals. But the model does not reflect evidence of declining discount rates over time, since in the model, the rate of discounting from one period to the next is constant after the first period.

The advantage of the quasi-hyperbolic discount function is that, like the DU model, it generates an infinite series that converges. This permits the calculation of the discounted present value of the sum of the infinite series:

$$1 + \sum_{t=1}^{T} \beta\delta^t = 1 + \beta\delta + \beta\delta^2 + \beta\delta^3 + \ldots = 1 + \beta\delta\left(\frac{1}{1-\delta}\right).$$

As before, the value of infinitely repeated payoffs can be factored out of the infinite series and then multiplied by the above expression to find the discounted present value of the series. Thus, the quasi-hyperbolic discount function shares the mathematical tractability of exponential discounting and the descriptive accuracy of hyperbolic discounting, particularly over the short term.

The quasi-hyperbolic discount function is presented along with the exponential and generalized hyperbolic discount functions in Figure 3. Here $\delta = .9$ for the exponential discount function; $\alpha = 2$ and $\gamma = 1$ for the hyperbolic discount function; and $\beta = .7$ and $\delta = .9$ for the quasi-hyperbolic discount function.

Applications to Psychology

Much of the inspiration for the research on discounting and intertemporal choice derives from psychology, where behavior relating to intrapersonal conflict, self-control, and addiction all point to problems with the standard DU model. Here, we briefly examine research on addictive behavior and strategies for self-control, which exhibit preference reversals that are difficult to reconcile with rationality assumptions.

We have all set the alarm clock for an early wake-up, reflecting a preference for getting up over sleeping in. Yet when the alarm clock actually goes off, we sometimes groggily press the snooze, reset the alarm, or simply turn it off, which reflects a preference reversal. We can think of countless other examples of situations in which an individual exhibits one set of preferences today but reverses himself the next day because of weakness or lack of self-control (Ainslie and Haslam 1992; Baumeister and Vohs 2003; O'Donoghue and Rabin 1999, 2003; Prelec and Bodner 2003). Scholars in a variety of disciplines—including political science, economics, and psychology—have conducted research on self-control. One finding is that the degree of self-control varies enormously over individuals and is quite consequential.

Many scholars have attempted to analyze problems of self-control in terms of a battle between multiple selves (Elster 1979; Winston 1980; Schelling 1984). Shefrin and Thaler (1992, 290) consider a multiple-selves model in which one self is a farsighted present-period decision maker, while the other is a myopic future consumer. There is a conflict in the intertemporal planning between the planner the night before and the "doer" of the next morning. Thus, the planner sets the alarm, and the doer resets it the next morning. Other researchers use models of self-control to study addiction (Winston 1980; Bickel and Johnson 2003) and other health-related issues, such as exercising and dieting (Chapman 2003; Herman and Polivy 2003).[18]

The standard solution to problems of self-control and addiction involve some kind of commitment strategy. The aim is to find a way for individuals to commit or bind themselves beforehand to a pattern of behavior, so that they either cannot go back on their commitment or find it very costly to do so (Elster 1979, Schelling 1992, Ainslie and Haslam 1992). Examples of such self-commitment devices include savings

Figure 3
Exponential, Hyperbolic, and Quasi-Hyperbolic Discount Functions
Exponential: $D(t) = \delta^t$, $\delta = .9$
Hyperbolic: $D(t) = (1 + \alpha t)^{-\gamma/\alpha}$, with $\alpha = 2$, $\gamma = 1$
Quasi-Hyperbolic: $D(t)$: $\{1, \beta\delta, \beta\delta^2, \beta\delta^3, \ldots\}$, $\delta = .9$, $\beta = .7$

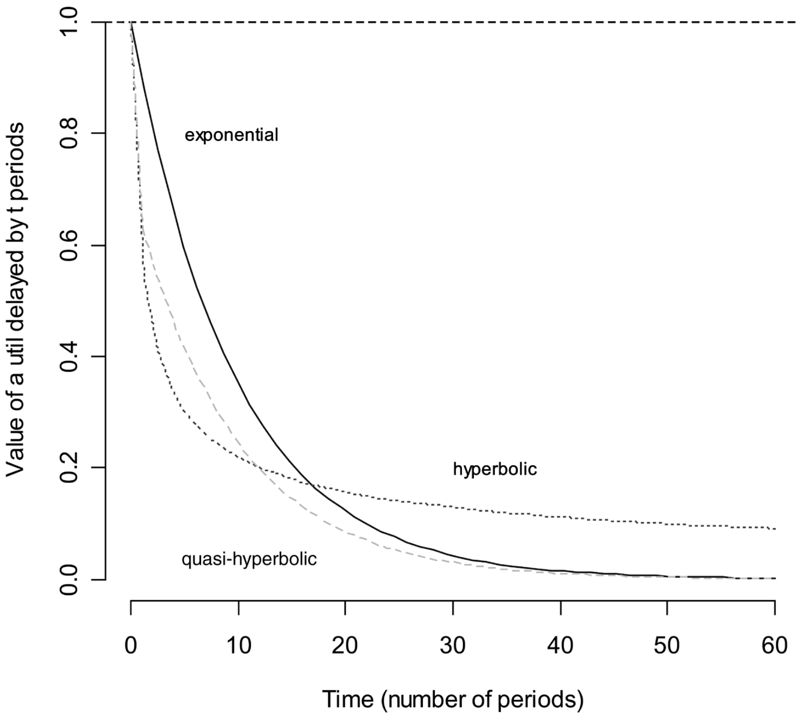

account withdrawal penalties or "Christmas Clubs" in banking or participating in clinics for alcoholics (O'Donoghue and Rabin 1999, 106-7). An often-cited example from addiction research is the drug Antabuse, the sole purpose of which is to make those who take it violently ill if they consume alcohol (Ainslie 1988).

The classic example of a commitment strategy comes from Homer's myth of Ulysses and the Sirens in the *Odyssey*. Ulysses learns from the goddess Circe that his ship will be traveling past the island of the Sirens, whose beautiful songs lure sailors in their direction and to their destruction on the rocks. Ulysses longs to hear the ravishing voices of the Sirens without suffering the disastrous consequences, and with that goal in mind, follows Circe's advice in developing a commitment strategy. He orders his crew of sailors to stuff their ears with wax, so that they cannot hear the Sirens, and to tie him to the mast of the ship. Anticipating his own preference reversal and his subsequent demands to be released so he can follow the call of the Sirens, Ulysses orders his crew to ignore any pleas he might make to be released from the mast until they are safely past the island (Elster 1979).

A number of social psychologists have dealt with issues of time horizons in research on "time perspective" (TP). TP research examines differences in individuals in temporal biases toward the past, present, or future (Zimbardo and Boyd 1999, 1272). A subset of TP dealing just with the relative bias of the future to the present, or future orientation, is called future time perspective (FTP), which Strathman et al. (1994, 742) describe as "a general preoccupation with the future or future events." An individual with less future orientation basically has a higher discount rate and is relatively myopic. An individual with greater future orientation discounts the future less—she is more forward-looking, more patient, and less myopic. Conceptually, FTP is similar to discounting in intertemporal choice. Both show how important the future is to the present, although the intertemporal choice literature is more rigorous in its specification of a discount rate and its consequences.[19]

Given these similarities, the literatures on intertemporal choice and on TP examine some similar questions, including drug dependence. Research on discounting shows that drug-dependent individuals discount the future greater than do non-drug-dependent individuals (Bickel and Johnson 2003, 435-6), while TP research shows that drug-dependent individuals are more present oriented (Keough, Zimbardo, and Boyd 1999), have a truncated FTP measure (Alvos, Gregson, and Ross 1993), or have "shortened time horizons" (Petry, Bickel, and Arnett 1998).

Applications to Political Science

While social psychologists and behavioral economists have devoted considerable attention to the ways in which people discount the future and make choices over time, political scientists have given relatively little attention to this issue. Some adopt models of constant-rate exponential discounting from economics, but they generally do so without questioning whether such an exponential discounting function adequately captures the way in which people, organizations, or states make choices over time. In this section, we attempt to demonstrate that the substitution of more descriptively accurate discounting functions, the quasi-hyperbolic discount function in particular, for the standard DU model can have significant consequences for models of strategic interaction in international relations.

We focus in particular on Robert Axelrod's (1984) influential analysis of cooperation in iterated Prisoner's Dilemma games. The single-play Prisoner's Dilemma game, with its dominant strategy of defection and equilibrium outcome of mutual conflict, has long been taken as a model of and metaphor for international conflict in an anarchic system with no higher authority to enforce agreements. In this context, Axelrod's (1984) demonstration—analytically and through a computer tournament—that under certain conditions, cooperation can emerge among egoistic actors in iterated Prisoner's Dilemma games, has had an enormous impact, not only in political science but in other fields as well. The implication is that international anarchy does not condemn states to mutual conflict and that strategies of reciprocity (and Tit-for-Tat [TFT] in particular) can lead to cooperation among self-interested actors who repeatedly interact in anarchic environments, provided that the shadow of the future is sufficiently long.[20]

Axelrod's analysis is based on the standard exponential discounting model from economics and his shadow of the future is the discount factor.[21] Our question is, what happens to cooperation if a more descriptively accurate discounting model is substituted for the exponential DU model in the iterated Prisoner's Dilemma game?[22]

In repeated plays of Prisoner's Dilemma, a strategy for each player specifies her actions (Cooperate or Defect) for each possible move in each future round of the game. For example, All-Defect (All-D), refers to a strategy in which a player always plays Defect, regardless of what her opponent did in the previous round. In TFT, a player cooperates on the first move and then in each round plays the same move that her opponent played in the previous round. If column player plays Cooperate in round t, then row player plays Cooperate in round $t + 1$.

In any given round, a player maximizes her payoff if she defects while her opponent cooperates. That will quickly generate a pattern in which both sides defect and in each subsequent round receive a lesser payoff associated with mutual defection. If a player begins by cooperating, however, she receives a lesser payoff in that round but induces her TFT opponent to cooperate on the next move and thereafter, providing each with a payoff superior to that of mutual defection, even if inferior to that for unilateral defection. The question is whether and at what point the benefits of cooperation over the long term outweigh the benefits of unilateral defection in the present. Axelrod (1984) demonstrates analytically that long-run mutual cooperation is possible if and only if the shadow of the future is sufficiently long, as indicated by a discount factor that is sufficiently high.[23]

Axelrod (1984) argues that while it is not possible to demonstrate conclusively that any one strategy is always better than all other strategies against the full range of strategies an opponent might play, it is possible to show that one strategy is better than another strategy against an adversary playing a particular strategy. Axelrod examines the conditions under which a TFT strategy will bring higher payoffs than an All-D strategy in a game with an adversary playing a TFT strategy.[24] This leads to the following inequality for a two-player iterated Prisoner's Dilemma game:

$$\text{Value of TFT v. TFT} \;\geq\; \text{Value of All-D v. TFT.}$$

Figure 4
Axelrod's (1984) Prisoner's Dilemma Game

	Cooperate	Defect
Cooperate	$R,R = 3, 3$ Reward for Mutual Cooperation	$S,T = 0, 5$ Sucker's Payoff; Temptation to Defect
Defect	$T,S = 5, 0$ Temptation to Defect; Sucker's Payoff	$P,P = 1, 1$ Reward for Mutual Defection

Using Axelrod's (1984) terminology for payoffs from Figure 4 (and the notation δ rather than *w* for the discount factor), we have

$$R + \delta R + \delta^2 R + \delta^3 R + \ldots \geq T + \delta P + \delta^2 P + \delta^3 P + \ldots$$

Given

$$1 + \delta + \delta^2 + \delta^3 + \ldots = \sum_{t=0}^{\infty} \delta^t = \frac{1}{1-\delta},$$

the inequality simplifies to

$$\delta \geq \frac{T-R}{T-P}.$$

If this inequality holds, δ is sufficiently high and therefore the shadow of the future is sufficiently long that TFT will dominate All-D against a player playing TFT, and cooperation can emerge.

It is important to emphasize that this analytically derived and theoretically consequential finding is based on the assumption that players discount the future at a constant rate, which generates an exponential discounting function. As we have argued, however, a growing body of experimental research on discounting behavior contradicts this assumption.

What happens in the same iterated Prisoner's Dilemma game if both players engage in the kind of discounting behavior commonly observed in experimental studies? We use a quasi-hyperbolic discount function rather than a hyperbolic function, which is not mathematically tractable and which does not permit an analytic solution.

Substituting the quasi-hyperbolic discount function $\{1, \beta\delta, \beta\delta^2, \beta\delta^3, \ldots\}$ into Axelrod's (1984) inequality,

$$R + \beta\delta R + \beta\delta^2 R + \beta\delta^3 R + \ldots \geq T + \beta\delta P + \beta\delta^2 P + \beta\delta^3 P + \ldots.$$

Simplifying, we have

$$R[1 + \beta\delta(1 + \delta + \delta^2 + \ldots)] \geq T + \beta\delta P(1 + \delta + \delta^2 + \ldots)$$

$$R\left(\frac{\beta\delta}{1-\delta}\right) - P\left(\frac{\beta\delta}{1-\delta}\right) \geq T - R$$

$$\beta\delta(R - P) \geq (1 - \delta)(T - R)$$

$$\beta\delta(R - P) + \delta(T - R) \geq T - R$$

$$\delta \geq \frac{T - R}{\beta(R - P) + T - R}$$

This expression specifies the conditions under which a TFT strategy will lead to higher payoffs than an All-D strategy against an opponent playing TFT if players engage in quasi-hyperbolic discounting. These conditions are a function of the payoffs of the basic game and β, which reduces by a constant proportion the weight of payoffs in every period beyond the present.

Note that the expression for δ is decreasing in β. When $\beta = 1$, the quasi-hyperbolic discount function reduces to the exponential discount function of the DU model, and $\delta \geq$ (T-R)/(T-P), the same result as in Axelrod (1984). As β decreases and moves away from 1, the threshold of δ required for cooperation increases. This can be inferred from Figure 3, which shows that at every point beyond the present, the value of the quasi-hyperbolic discount function is less than that of the exponential discount function for a given δ.

Thus, with quasi-hyperbolic discounting and with $\beta < 1$, the threshold that the discount factor must meet for TFT to dominate All-D is greater than under standard exponential discounting. It takes a higher δ or shadow of the future to generate an equivalent discounted present value and hence defeat an All-D strategy. This means that if players' time preferences are similar to those repeatedly observed in countless laboratory studies regarding a range of intertemporal choice problems, so that their time preferences are better modeled by a quasi-hyperbolic discount function than by an exponential discount function, then cooperation in iterated Prisoner's Dilemma games is more difficult than Axelrod (1984) suggests. Stated differently, to the extent that people behave as quasi-hyperbolic discounters rather than exponential discounters, the conditions for cooperation in an iterated Prisoner's Dilemma game are somewhat more restrictive than Axelrod's (1984) analysis suggests.

A word of caution is in order. We must remember that the quasi-hyperbolic discount function does not fully capture all of the experimentally observed deviations from the standard DU model. It reflects the heavy discounting of the near future. It can also account for the common difference effect and the preference reversals associated with them, given the steep discounting of the near future. But since the quasi-hyperbolic discounter engages in constant-rate discounting after the period after the present, it does not entirely reflect the experimentally observed pattern of declining discount rates, which result in higher weights attached to the future than predicted by the exponential model.[25]

Our results are theoretically valid for players who engage in quasi-hyperbolic discounting. They raise the question, however, of what happens to cooperation if players not only heavily discount the near future but also if they apply declining discount rates to the more distant future and thus give payoffs from more distant rounds of the game greater weight in their current calculations than would an exponential discounter. At some point, the cumulation of those payoffs might counteract the effect of the steep discounting of the period after the present and increase the relative benefits of cooperation.

This would be a useful project for future research, but it would quickly get complicated because of the lack of convergence of the hyperbolic discount function. Analysts might use computational modeling to estimate the discounted present values of various strategies with different parameter values. Such an analysis would also have to involve a specification of how many periods into the future individuals incorporate into their present calculations. That is an important question, and it is relevant for exponential discounting as well as for hyperbolic discounting, but it goes beyond the scope of the current study.[26]

Summary and Conclusions

We have summarized the DU model, which incorporates constant-rate discounting and an exponential function and which has been standard in economics and related fields since Samuelson (1937) introduced it seven decades ago. We have identified some of the leading behavioral anomalies in the DU model, as discovered by experimental economists and social psychologists. These anomalies include declining discount rates over time and hence the continually increasing relative weight given to future payoffs; the common difference effect and the preference reversals it generates, given that an actor may prefer x now to y in the next period but prefer y in $t+1$ periods from now instead of x in t periods from now; the absolute magnitude effect, or the tendency for people to discount smaller payoffs at a higher rate than larger payoffs; the gain-loss asymmetry, or the tendency for people to discount gains at a higher rate than they discount losses (which, consistent with the spirit of prospect theory, means that future losses have a greater impact than do future gains); the delay-speedup asymmetry, in which the psychological cost of a delay in receiving an expected reward is greater than the psychological benefit of receiving the same reward sooner than expected; and preferences for improving sequences, in which people on average prefer sequences of payoffs with improving values to sequences of declining values.

Each of these anomalies represents a violation of a key assumption of the DU model. The first two anomalies bear on the DU model's assumption of constant-rate discounting. The third and fourth anomalies bear on the DU model's assumptions that discounting is independent of the utility of the payoffs. The fourth and fifth anomalies reveal the importance of reference points and framing, independent of the formal mathematical structure of a choice problem. The sixth anomaly contradicts the assumption of positive time preference.

After summarizing the leading behavioral anomalies in the DU model and considering some of their consequences, we examined some of the leading alternative models that have been constructed to account for some of these anomalies. We gave particular attention to the hyperbolic model, which captures the phenomena of declining discount rates, common difference effect, and preference reversals but which cannot account for the other observed anomalies in the DU model. We also discussed the quasi-hyperbolic model, which, like the hyperbolic model, captures the steep discounting of the near future, the common difference effect, and its associated preference reversals but which does not capture declining discount rates in the more distant future. We emphasized that the lack of convergence of the hyperbolic discount function has generally limited its application to an analysis of the discounted present values of a handful of discrete future outcomes, not a repeated series of outcomes, while the quasi-hyperbolic discount function, like the exponential function, can be easily applied to repeated series. Thus, in many respects, the quasi-hyperbolic discount model provides a middle ground between the exponential discount model and hyperbolic discount model.

We then applied the quasi-hyperbolic discount model to the question of the conditions for cooperation in iterated Prisoner's Dilemma games. We demonstrated that if actors in such a game reflect the widely observed tendency of individuals in other settings to discount the immediate future at a higher rate than the more distant future so that their behavior is better modeled by a quasi-hyperbolic discount function than by a constant-rate exponential function, then the conditions under which cooperation will emerge in an iterated Prisoner's Dilemma game are somewhat more restrictive than Axelrod (1984) suggests.

We also suggested some possible directions for future research. One involves the analysis of the conditions for cooperation in an iterated Prisoner's Dilemma game among hyperbolic (as opposed to quasi-hyperbolic) discounters. This might require computational modeling, and it would have to incorporate the number of periods an individual looks into the future.

Another interesting direction for future research would be an examination of the effects of the magnitude of the payoffs and the asymmetry of losses and gains on cooperation in iterated Prisoner's Dilemma games. Axelrod (1984) specifies the payoffs of each stage of his iterated Prisoner's Dilemma game as DC = 5, CC = 3, DD = 1, and CD = 0, where D is a strategy of Defect and C is a strategy of Cooperation. The implication is that nothing would change if the payoffs were multiplied or divided by a constant amount, since players' preference orders and their rations would be unchanged. The research on anomalies leads us to question this prediction. If all payoffs were multiplied by a factor greater than one, the tendency for people to apply lower discount rates to larger payoffs would result in greater weight to the future, which would make cooperation easier. Similarly, if all payoffs were multiplied by a factor less than one, then we would predict that the relatively higher discount rates applied to smaller payoffs would lower the weight given to the future and make cooperation more difficult. The question is by how much. What is the impact of a marginal and proportional increase in payoffs on the likelihood of cooperation?

We could apply similar logic to negative payoffs, assuming reference-dependent utilities, and explore the impact of the differential impact of losses and gains on discounting and consequently on the likelihood of cooperation, given the tendency for people to discount gains at a higher rate than losses and thus give future losses greater weight than future gains. This would be somewhat more complicated, since the asymmetry of losses and gains might affect cooperation through other causal paths, which the analyst would have to separate. Individuals making choices among losses tend to be more risk acceptant than those making choices among gains (Kahneman and Tversky 1979), but risk orientation has no impact in a complete information game. More relevant are the direct effects of loss aversion or the tendency to give losses (present as well as future) greater weight than comparable gains. The analyst would need to separate those effects from the indirect effects of negative payoffs on outcomes through their impact on individual discount rates.

More generally, we need to think more about the time horizons of political leaders, how time horizons vary over individuals and regime types and cultures, and with what consequence for foreign policy decisions and strategic interaction.

NOTES

1. One could include a different discount rate for different goods, although this would complicate the analysis, particularly tradeoffs between goods.

2. In Axelrod's (1984) iterated Prisoner's Dilemma game, for example, a high discount factor w is associated with a longer "shadow of the future."

3. A good example of a discount factor approximating zero is the comment of a member of a tribe of hunter-gatherers who had previously had no contact with modern civilization and who left the Columbian jungle to join the modern world. He said, "The future, what's that?" (Forero 2006).

4. The last, while acknowledged, is rarely included in standard lists of behavioral anomalies in the model.

5. The corresponding discount factors are $\delta = .66$, .72, and .79, respectively.

6. Preference reversals are sufficiently common that Henderson and Langford (1998, 1494) describe them as "a normal part of human behavior."

7. Preference reversals can also occur for two sets of more distant outcomes, but the effect would be less dramatic since discount rates decline over time.

8. The corresponding discount factors are $\delta = .25$, .71, and .75, respectively.

9. The corresponding discount factors are $\delta = .48$ and .66, respectively.

10. For an alternative theory of reference-dependent preferences, see Munro and Sugden (2003). For applications of prospect theory to international relations and to political science, see Levy (2000, 2003).

11. As Camerer, Loewenstein, and Prelec (2005, 10) note, advances in neuroscience allow us to move past revealed preference theory in the study and measurement of utility.

12. We would also predict that the magnitude of losses in declining sequences will increase while the magnitude of gains in ascending sequences will decline, given the tendency (the "instant endowment effect") for people to "renormalize" their reference point after gains but not after losses (Kahneman and Tversky 1979).

13. Most applications of hyperbolic discounting models in psychology and in behavioral economics have focused on issues of self-control and commitment, which we discuss in a subsequent section.

14. This suggests the possibility that discounting has biological origins.

15. As α approaches zero, the hyperbolic discount function approximates the exponential function, while very high values of α make the function look like a step function with a near-vertical decline in the first period and a flat near-horizontal curve thereafter (Loewenstein and Prelec 1992, 580).

16. Loewenstein and Prelec (1992) propose a model to account for each of these anomalies except the preference for improving sequences. The model includes both the generalized hyperbolic discount function and a two-piece value function that incorporates a reference point and two separate functions, one to represent gains and the other to represent losses. Thus far, however, few others have adopted this rather complex model.

17. Another reason hyperbolic discounting functions are not applied to iterated series relates to the concern that the independence of utility assumption, which is needed for the separability of utility to sum the utilities into the intertemporal utility function, is "far from innocuous" (Frederick, Loewenstein, and O'Donoghue 2003, 21). If individuals can have preferences over entire sequences of utilities, as the anomaly of improving sequences suggests, then this assumption is problematic (Loewenstein and Prelec 1992, 578).

18. Strotz (1956), Pollak (1968), and O'Donoghue and Rabin (1999, 2003) take a different approach. Instead of using a model based on dual personalities, they distinguish between "sophisticates" and "naifs." Sophisticates are those who understand that their preferences will change in the future and that they will encounter problems of self-control, while naifs do not. As a result, naifs will procrastinate when costs are immediate and act immediately when the rewards are immediate. Sophisticates, on the other hand, always carry out decisions immediately, regardless of whether rewards or costs are immediate, because they know their preferences can change in the future.

19. Another topic in psychology that focuses on time and utility is "affective forecasting," which examines how well people anticipate the psychological value they derive from future outcomes (Gilbert and Ebert 2002; Kahneman, Diener, and Schwarz 2003).

20. A long shadow of the future has other implications. Fearon (1998), who conceives of cooperation as involving an initial bargaining stage and subsequent enforcement stage, argues that a long shadow of the future increases the importance of the outcomes of earlier iterations of the game, which will define the parameters of later iterations of the game. If the shadow of the future is long, actors will bargain harder in early stages, and cooperation will be more difficult. This is useful, but it does not bear directly on the question of cooperation in iterated Prisoner's Dilemma games, as defined by Axelrod (1984), where there is no bargaining and where the parameters of each stage of the game are identical. It may be true that this assumption is an unrealistic one for

many real world problems, particularly security issues, and that where such bargaining occurs, Fearon's (1998) model may be useful. But it is technically a different game than the one proposed by Axelrod (1984).

Blaydes (2004) provides an interesting addition to Fearon's argument: She argues that small states with lesser discount factors can have greater bargaining strength by threatening to defect. Cooperation is sustained through an increase in rewards for the impatient state. Empirically, she operationalizes the discount factor by arguing that Organization of Petroleum Exporting Countries members that are poor and more populous or have unstable regimes have greater impatience and thus lesser discount factors.

21. Many game theorists have praised Axelrod (1984) for drawing popular attention to evolutionary behavior but have argued that too many scholars have read too much into Axelrod's findings. Game theorists emphasize that Axelrod's basic finding is not new but in fact was demonstrated in folk theorems that scholars had advanced by the 1950s (Binmore 1992, 369-77; Gintis 2000, 126-29; Morrow 1994, 268-79). Second, although Axelrod (1984, 1997) acknowledged that the utility of the strategy of TFT is a function of the strategies adopted by others, many scholars have erroneously come to regard TFT as a strategy that is optimal under a wide variety of circumstances. Linster (1992) shows that a strategy of Grim Trigger (cooperate until the adversary defects and then defect forever) outperforms TFT in many situations, including those in which the evolutionary process is subjected to constant perturbations. Game theorists also criticize certain aspects of Axelrod's computer tournament. Axelrod's simulation involved pairwise interactions, whereas in a game of multiactor interactions, anyone can punish a defector (Binmore 1998). Axelrod's simulation involved a finite number of periods, but none of the strategies entered in the tournament were programmed to exploit the end effects of a finitely repeated game (Binmore 1998). Finally, theorists have questioned Axelrod's definition of evolutionary stability and his argument (Axelrod and Hamilton 1981) that TFT is an evolutionarily stable strategy (Boyd and Lorberbaum 1987; Bendor and Swistak 1997).

22. Cooperation is also easier if people are not purely self-interested but instead altruistic or driven by the logic of obligatory action or norms of social appropriateness (March and Olsen 1998, 951; Goldgeier and Tetlock 2001, 82). This probably explains why we observe higher than expected levels of cooperation in experimental situations. That cooperation can emerge even among self-interested egoists is the more striking finding, however, and deserves special attention.

With respect to altruism, it is important to distinguish between genuine altruism and strategic incentives to build a reputation for altruism in early stages of the game if there is some chance that the other player is altruistic, as Andreoni and Miller (1993) show in their work on the sequential equilibrium reputation model. On norms, see also Axelrod (1997).

23. Axelrod's (1984) computer tournament assumed perfect information of each player's previous moves. However, subsequent research shows that uncertainty as to whether the adversary has played Cooperate or Defect can impede cooperation under Tit-for-Tat (TFT) and thus increase the effectiveness of more forgiving strategies that allow for the possibility of chance defections (Bendor 1993).

24. There are a multitude of strategies that could be compared with TFT in place of All-Defect (All-D), but we stay with the comparison between TFT and All-D to stay consistent with Axelrod (1984).

25. Nor does it reflect framing effects, but those are not relevant to the standard version of Axelrod's (1984) Prisoner's Dilemma game with reference-independent payoffs.

26. Thus, time horizons are best conceptualized as a multidimensional concept that includes both the discount rate and a degree of myopia or how many periods one looks into the future.

References

Ainslie, G. 1974. Impulse control in pigeons. *Journal of the Experimental Analysis of Behavior* 21:485-89.

———. 1975. Specious reward: A behavioral theory of impulsiveness and impulse control. *Psychological Bulletin* 82:463-96.

———.1988. Matching is the integrating framework. *Behavioral and Brain Science* 11:679-80.

Ainslie, G., and V. Haendel. 1983. The motives of the will. In *Etiologic aspects of alcohol and drug abuse*, edited by E. Gottheil, K. Duley, T. Skoloda, and M. Waxman, 119-40. Springfield, IL: Charles C Thomas.

Ainslie, G., and N. Haslam. 1992. Self-control. In *Choice over time*, edited by G. Loewenstein and J. Elster, 177-209. New York: Russell Sage.

Alvos, L., R. Gregson, and M. Ross. 1993. Future time perspectives in current and previous injecting drug users. *Drug and Alcohol Dependence* 31:193-97.

Andreoni, J., and J. Miller. 1993. Rational cooperation in the finitely repeated Prisoner's Dilemma: Experimental evidence. *The Economic Journal* 103:570-85.

Angeletos, G., D. Laibson, J. Tobacman, A. Repetto, and S. Weinberg. 2003. The hyperbolic consumption model: Calibration, simulation, and empirical evaluation. In *Time and decision,* edited by G. Loewenstein, D. Read, and R. Baumeister, 517-543. New York: Russell Sage.

Axelrod, R. 1984. *The evolution of cooperation*. New York: Basic Books.

———. 1997. *The complexity of cooperation*. Princeton, NJ: Princeton University Press.

Axelrod, R., and W. Hamilton. 1981. The evolution of cooperation. *Science* 211:1390-96.

Barro, R. 1999. Ramsey meets Laibson in the neoclassical growth model. *The Quarterly Journal of Economics* 114:1125-52.

Baumeister, R., and K. Vohs. 2003. Willpower, choice, and self-control. In *Time and decision,* edited by G. Loewenstein, D. Read, and R. Baumeister, 201-16. New York: Russell Sage.

Bendor, J. 1993. Uncertainty and the evolution of cooperation. *Journal of Conflict Resolution* 37:709-34.

Bendor, J., and P. Swistak. 1997. The evolutionary stability of cooperation. *American Political Science Review* 91:290-307.

Benzion, U., A. Rapoport, and J. Yagil. 1989. Discount rates inferred from decisions: An experimental study. *Management Science* 35:270-84.

Berns, G., J. Chappelow, M. Cekic, C. Zink, G. Pagnoni, and M. Martin-Skurski. 2006. Neurobiological substrates of dread. *Science* 312:754-58.

Bickel, W., and M. Johnson. 2003. Delay discounting: A fundamental process of drug dependence. In *Time and decision,* edited by G. Loewenstein, D. Read, and R. Baumeister, 419-40. New York: Russell Sage.

Binmore, K. 1992. *Fun and games*. Lexington, MA: D.C. Heath.

———. 1998. Review of *The complexity of cooperation* by R. Axelrod. *Journal of Artificial Societies and Social Simulation* 1. http://jasss.soc.surrey.ac.uk/1/1/review1.html (accessed on October 14, 2006).

Blaydes, L. 2004. Rewarding impatience: A bargaining and enforcement model of OPEC. *International Organization* 58:213-37.

Boyd, R., and J. Lorberbaum. 1987. No pure strategy is evolutionarily stable in the repeated Prisoner's Dilemma game. *Nature* 327:58-59.

Camerer, C., G. Loewenstein, and D. Prelec. 2005. Neuroeconomics: How neuroscience can inform economics. *Journal of Economic Literature* 43:9-64.

Chapman, G. 1996. Temporal discounting and utility for health and money. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22:771-91.

———. 2003. Time discounting of health outcomes. In *Time and decision,* edited by G. Loewenstein, D. Read, and R. Baumeister, 395-417. New York: Russell Sage.

Chung, S., and R. Herrnstein. 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.

Elster, J. 1979. *Ulysses and the sirens*. Cambridge, UK: Cambridge University Press.

Elster, J., and G. Loewenstein. 1992. Utility from memory and anticipation. In *Choice over time*, edited by G. Loewenstein and J. Elster, 213-34. New York: Russell Sage.

Fearon, J. 1998. Bargaining, enforcement, and international cooperation. *International Organization* 52:269-306.

Forero, J. 2006. Leaving the wild, and rather liking the change. *New York Times*, May 11, A1.

Frederick, S. 2003. Time preference and personal identity. In *Time and decision,* edited by G. Loewenstein, D. Read, and R. Baumeister, 89-113. New York: Russell Sage.

Frederick, S., G. Loewenstein, and T. O'Donoghue. 2003. Time discounting and time preference. In *Time and decision,* edited by G. Loewenstein, D. Read, and R. Baumeister, 13-86. New York: Russell Sage.

Ganiats, T., R. Carson, R. Hamm, S. Cantor, W. Sumner, S. Spann, M. Hagen, and C. Miller. 2000. Health status and preferences: Population-based time preferences for future health outcomes. *Medical Decision Making* 20:263-70.

Gilbert, D., and J. Ebert. 2002. Decisions and revisions: The affective forecasting of changeable outcomes. *Journal of Personality and Social Psychology* 82:503-14.

Gintis, H. 2000. *Game theory evolving*. Princeton, NJ: Princeton University Press.

Goldgeier, J., and P. Tetlock. 2001. Psychology and international relations theory. *Annual Review of Political Science* 4:67-92.

Green, L., A. Fry, and J. Myerson. 1994. Discounting of delayed rewards: A life-span comparison. *Psychological Science* 5:33-36.

Henderson, N., and I. Langford. 1998. Cross-disciplinary evidence for hyperbolic social discount rates. *Management Science* 44:1493-1500.

Herman, P., and J. Polivy. 2003. Dieting as an exercise in behavioral economics. In *Time and decision,* edited by G. Loewenstein, D. Read, and R. Baumeister, 459-89. New York: Russell Sage.

Herrnstein, R. 1981. Self-control as response strength. In *Quantification of steady-state operant behavior*, edited by C. Bradshaw, E. Szabadi, and C. Lowe, 3-20. New York: Elsevier-North Holland.

———. 1990. Rational choice: Necessary but not sufficient. *American Psychologist* 45:356-67.

Hsee, C., R. Abelson, and P. Salovey. 1991. The relative weighting of position and velocity in satisfaction. *Psychological Science* 2:263-66.

Kahneman, D., E. Diener, and N. Schwarz, eds. 2003. *Well-being: The foundations of hedonic psychology*. New York: Russell Sage.

Kahneman, D., and A. Tversky. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* 47:363-91.

Keough, K., P. Zimbardo, and J. Boyd. 1999. Who's smoking, drinking, and using drugs? Time perspective as a predictor of substance use. *Basic and Applied Social Psychology* 21:149-64.

Kirby, K. 1997. Bidding on the future: Evidence against normative discounting of delayed rewards. *Journal of Experimental Psychology* 126:54-70.

Kirby, K., and R. Herrnstein. 1995. Preference reversals due to myopic discounting of delayed reward. *Psychological Science* 6:83-89.

Kirby, K., and N. Marakovic. 1995. Modeling myopic decisions: Evidence for hyperbolic delay-discounting with subjects and amounts. *Organizational Behavior and Human Decision Processes* 64:22-30.

———. 1996. Delay-discounting probabilistic rewards: Rates decrease as amounts increase. *Psychonomic Bulletin and Review* 3:100-04.

Laibson, D. 1997. Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics* 112: 443-77.

———. 2001. A cue-theory of consumption. *Quarterly Journal of Economics* 116:81-119.

Laibson, D., A. Repetto, J. Tobacman, R. Hall, W. Gale, and G. Akerlof. 1998. Self-control and saving for retirement. *Brookings Papers on Economic Activity* 1:91-196.

Levy, J. S. 2000. The implications of framing and loss aversion for international conflict. In *Handbook of war studies II*, edited by M. Midlarsky, 193-221. Ann Arbor: University of Michigan Press.

———. 2003. Applications of prospect theory to political science. *Synthese* 135:215-41.

———. 2007. Preventive war and the Bush doctrine: Theoretical logic and historical roots. In *The Bush Doctrine: Psychology and strategy in an age of terrorism*, edited by S. Renshon and P. Suedfeld, 175-200. London: Routledge.

Linster, B. 1992. Evolutionary stability in the infinitely repeated Prisoners' Dilemma played by two-state Moore machines. *Southern Economic Journal* 58:880-903.

Loewenstein, G. 1987. Anticipation and the valuation of delayed consumption. *The Economic Journal* 97:666-84.

———. 1988. Frames of mind in intertemporal choice. *Management Science* 34:200-14.

———. 1992. The fall and rise of psychological explanations in the economics of intertemporal choice. In *Choice over time*, edited by G. Loewenstein and J. Elster, 3-34. New York: Russell Sage.

———. 2006. The pleasures and pains of information. *Science* 312:704-06.

Loewenstein, G., and D. Prelec. 1992. Anomalies in intertemporal choice: Evidence and an interpretation. *The Quarterly Journal of Economics* 107:573-97.

———. 1993. Preferences for sequences of outcomes. *Psychological Review* 100:91-108.

Loewenstein, G., and N. Sicherman. 1991. Do workers prefer increasing wage profiles? *Journal of Labor Economics* 9:67-84.

March, J., and J. Olsen. 1998. The institutional dynamics of international political orders. *International Organization* 52:943-69.

Mazur, J. 1987. An adjustment procedure for studying delayed reinforcement. In *The effect of delay and intervening events on reinforcement value,* edited by M. Commons, J. Mazur, J. Nevin, and H. Raiffa, 55-73. Hillsdale, NJ: Lawrence Erlbaum.

Morrow, James. 1994. *Game theory for political scientists*. Princeton, NJ: Princeton University Press.

Munro, A., and R. Sugden. 2003. On the theory of reference-dependent preferences. *Journal of Economic Behavior and Organization* 50:407-28.

O'Donoghue, T., and M. Rabin. 1999. Doing it now or later. *The American Economic Review* 89:103-24.

———. 2003. Self-awareness and self-control. In *Time and decision,* edited by G. Loewenstein, D. Read, and R. Baumeister, 217-43. New York: Russell Sage.

Pender, J. 1996. Discount rates and credit markets: Theory and evidence from rural India. *Journal of Development Economics* 50:257-96.

Petry, N., W. Bickel, and M. Arnett. 1998. Shortened time horizons and insensitivity to future consequences in heroin addicts. *Addiction* 93:729-38.

Phelps, E., and R. Pollak. 1968. On second-best national saving and game-equilibrium growth. *Review of Economic Studies* 35:185-99.

Pollak, R. 1968. Consistent planning. *Review of Economic Studies* 35:201-08.

Prelec, D. 2004. Decreasing impatience: A criterion for non-stationary time preference and "hyperbolic" discounting. *Scandinavian Journal of Economics* 106:511-32.

Prelec, D., and R. Bodner. 2003. Self-signaling and self-control. In *Time and decision,* edited by G. Loewenstein, D. Read, and R. Baumeister, 277-98. New York: Russell Sage.

Rachlin, H., and L. Green. 1972. Commitment, choice, and self-control. *Journal of the Experimental Analysis of Behavior* 17:15-22.

Rachlin, H., and A. Raineri. 1992. Irrationality, impulsiveness, and selfishness as discount reversal effects. In *Choice over time*, edited by G. Loewenstein and J. Elster, 93-118. New York: Russell Sage.

Samuelson, P. 1937. A note on measurement of utility. *Review of Economic Studies* 4:155-61.

———. 1939. The rate of interest under ideal conditions. *The Quarterly Journal of Economics* 53:286-97.

———. 1952. Probability, utility, and the independence axiom. *Econometrica* 20:670-78.

Schelling, T. 1984. Self-command in practice, in policy, and in a theory of rational choice. *American Economic Review* 74:1-11.

———. 1992. Self-command: A new discipline. In *Choice over time*, edited by G. Loewenstein and J. Elster, 167-76. New York: Russell Sage.

Shefrin, H., and R. Thaler. 1992. Mental accounting, saving, and self-control. In *Choice over time*, edited by G. Loewenstein and J. Elster, 287-330. New York: Russell Sage.

Strathman, A., F. Gleicher, D. Boninger, and C. Edwards. 1994. The consideration of future consequences: Weighing immediate and distant outcomes of behavior. *Journal of Personality and Social Psychology* 66:742-52.

Strotz, R. 1956. Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies* 23:165-80.

Thaler, R. 1981. Some empirical evidence on dynamic inconsistency. *Economics Letters* 8:201-7.

Winston, G. 1980. Addiction and backsliding. *Journal of Economic Behavior and Organization* 1:295-324.

Zimbardo, P., and J. Boyd. 1999. Putting time in perspective: A valid, reliable individual-differences metric. *Journal of Personality and Social Psychology* 77:1271-88.