Journal of Experimental Psychology:
Learning, Memory, and Cognition
2013, Vol. 39, No. 4, 1274–1279

0278-7393/13/$12.00 DOI: 10.1037/a0031170

RESEARCH REPORT

Discounting of Delayed Rewards Is Not Hyperbolic

Christian C. Luhmann
Stony Brook University

Delay discounting refers to decision-makers' tendency to value immediately available goods more than identical goods available only after some delay. In violation of standard economic theory, decision-makers frequently exhibit dynamic inconsistency; their preferences change simply due to the passage of time. The standard explanation for this behavior has appealed to the nature of decision-makers' discount functions, specifically positing a hyperbolic discount function. Though this explanation has been largely accepted, there has been surprisingly little work examining whether preference reversals are actually consistent with hyperbolic discounting. The current study holds hyperbolic discounting to the same empirical standard that exponential discounting has been held to and finds that choice behavior is not consistent with hyperbolic discounting. Despite the overwhelming focus placed on hyperbolic discounting, the current findings cast doubt on hyperbolic discounting as an explanation of decision-makers' undesirable preference reversals and as an explanation of delay discounting behavior in general.

Keywords: delay discounting, intertemporal choice, decision making, dynamic inconsistency, hyperbolic

Delay discounting refers to the robust finding that animals, including humans, behave as though immediately consumable goods are more valuable than those only available after some delay. For example, a decision-maker might choose $100 delivered immediately over $200 to be delivered in 3 years. One of the major empirical questions regarding such choices has focused on the discount function—the mathematical function specifying the relationship between reward magnitude, delay, and subjective value. Initial theoretical work in economics (Samuelson, 1937) suggested that discounting should be exponential, in which the subjective value of a delayed reward can be expressed as:

$$V_D = V_0 \cdot e^{-kD} \quad (1)$$

where V_D represents the current value of a reward that will be delivered after a delay D, k represents the decision-maker's discount rate, and V_0 represents the undiscounted value of that same reward (i.e., the value of that reward if it were available immediately). Unfortunately, the vast majority of empirical evidence has demonstrated that humans (Kirby, 1997; Kirby & Herrnstein, 1995) and non-human animals (Rachlin, 2006) do not discount exponentially. Herrnstein's matching law instead suggests that the subjective value of delayed rewards should be an inverse function of delay (Chung & Herrnstein, 1967), which has motivated what is now referred to as hyperbolic discounting, in which the subjective value of a delayed reward can be expressed according to Equation 2 with the same quantities defined as above.

$$V_D = \frac{V_0}{1 + kD} \quad (2)$$

The distinction between exponential and hyperbolic discounting may, at first, appear to be a bit of mathematical nitpicking. However, these two different discounting schemes represent very different conceptualizations of how delayed rewards are evaluated and have serious implications for behavior. Exponential discounting represents how bank loans work; the value of a delayed reward declines by a fixed percentage per unit of time. Hyperbolic discounters, on the other hand, behave as though their discount rates increase as the delivery of a delayed reward draws near. Thus, hyperbolic discounters appear to exhibit patience when dealing with rewards in the distant future only to find that this patience diminishes as rewards move closer in time.

The critical problem with the diminishing patience exhibited by hyperbolic discounters is that it tends to produce preferences that shift in systematic and predicable ways due to the simple passage of time; what economists refer to as dynamic inconsistency. For example, a hyperbolic discounter might prefer $200 delivered in 2 years over $100 delivered in 1 year (a patient preference) only to reverse this preference a year later, preferring an immediate $100 over $200 delivered in 1 year (an impatient preference).[1] Because the former pair of rewards will become the latter pair of rewards in 1 year, the preferences described by hyperbolic discounting are

This article was published Online First January 28, 2013.

I thank Elizabeth Altman and the Cognition and Decision-Making Lab for help with data collection and Michael Bixter for comments on a draft of this article.

Correspondence concerning this article should be addressed to Christian C. Luhmann, Department of Psychology, Stony Brook University, Stony Brook, NY 11794. E-mail: christian.luhmann@stonybrook.edu

[1] To be clear, hyperbolic discounting does not imply that preferences over an arbitrary pair of rewards will reverse for any change in delay. Whether or not a hyperbolic discounter reverses their preferences as time passes depends on the specific reward magnitudes, delays, and discount rates. That being said, changes in delay always alter the *strength* of a hyperbolic discounter's preference (i.e., the difference between the discounted value of the two rewards) in the predicted direction, even when actual choices do not reverse. In addition, reward pairs and an appropriate delay can always be found such that any non-exponential discounter (hyperbolic or otherwise) is guaranteed to reverse their preference (and thus represent an arbitrage opportunity).

contradictory. The predictable nature of these preference reversals means that such discounters can be exploited for essentially risk-free profits (i.e., arbitrage). For these reasons, hyperbolic discounting has been characterized as irrational (Ainslie, 2001; Kirby, 1997; Soman et al., 2005).

Unlike hyperbolic discounting, the preferences of exponential discounters do not change as a function of time, and thus preference reversals are simply not possible. This stability relies on the fact that exponential preferences are sensitive to the difference between the two delays but insensitive to the overall delay. This property of exponential discounting is referred to as the stationarity axiom (Koopmans, 1960) and guarantees that an exponential discounter will never exhibit dynamic inconsistency.

Empirical Evidence for Dynamic Inconsistency

Much of the empirical work on delay discounting has attempted to provide an increasingly precise estimate of the discounting function itself. This work typically estimates discount functions point by point and then fits exponential and hyperbolic functions to the result. These studies nearly always find a small advantage for hyperbolic functions, which, as explained above, implies that these decision-makers should exhibit dynamic inconsistency, at least in some circumstances. Though this is certainly a reasonable method for assessing preferences, this approach is unsatisfyingly indirect when one is specifically interested in dynamic inconsistency. No preference reversals are actually observed during the course of such an experiment. Instead, the reversals are simply inferred based on the best-fitting discount function. Surprisingly, there have been relatively few attempts to directly evaluate the preference reversals predicted under hyperbolic discounting.

The earliest empirical investigations of dynamic inconsistency involved non-human animals. Rachlin and Green (1972) developed a paradigm in which they could manipulate what has since been referred to as front-end delay. In this paradigm, decision-makers are offered a choice between a smaller reward available after a short delay (T) and a larger reward available after a longer delay ($T + 4$). Here, T refers to the front-end delay, and its manipulation allows experimenters to alter the absolute delay of a pair of rewards without altering the difference between the shorter and longer delays. Exponential discounting is only sensitive to the difference between the two delays (obeying the stationarity axiom); the front-end delay has no influence on the preferences of an exponential discounter. Thus, decision-makers that are sensitive to manipulations of front-end delay must be discounting delayed rewards non-exponentially. Hyperbolic discounting, in contrast, is sensitive to both the difference between the two delays and the front-end delay. As will be seen, this basic paradigm has been an extremely popular way of assessing stationarity.

Investigating the behavior of pigeons, Rachlin and Green (1972) found that, as the pair of rewards was shifted further into the future (i.e., as the front-end delay, T, was increased), patience appeared to increase with the larger, more delayed reward being chosen more frequently. As the pair of rewards was shifted to be closer in time (i.e., as the front-end delay, T, was decreased), patience appeared to decrease with the less delayed reward being chosen more frequently. Within the non-human animal literature, there have been several studies employing this same basic design (Ainslie & Herrnstein, 1981; Green & Estle, 2003; Green, Fisher, Perlow, & Sherman, 1981), and all have observed similar shifts in preference. Of these studies, only Ainslie and Herrnstein (1981) attempted to compare the observed preference reversals with those predicted by hyperbolic discounting (i.e., Equation 2). Though their evaluation was qualitative in nature, they did demonstrate that the preference reversals exhibited by their pigeons were at least approximately consistent with the predictions of hyperbolic discounting.

Evidence for dynamic inconsistency has also been observed in human decision-makers. Early studies (Millar & Navarick, 1984; Solnick, Kannenberg, Eckerman, & Waller, 1980) reported preference reversals using vastly simplified versions of the front-end delay procedure (e.g., only two front-end delays). More recent studies (Green, Fristoe, & Myerson, 1994; Kirby & Herrnstein, 1995) have employed more standard front-end delay paradigms (involving many front-end delays) allowing for more precise measurement of stationarity violations. In all experimental conditions, preferences were found to reverse as the front-end delay was increased.

The observation of stationarity violations reported in these studies is strong evidence against exponential discounting, but roughly consistent with hyperbolic discounting. However, despite the precision with which many studies can describe individuals' preferences, researchers have largely neglected to examine whether the observed reversals conform to the quantitative predictions of hyperbolic discounting. One exception to this trend is a study by Holt, Green, Myerson, and Estle (2008). In this study, similar methods were employed to evaluate the stationarity of preferences about delayed losses (rather than gains). After observing violations of stationarity, Holt et al. did fit their data with a hyperbolic discounting model. Though the resulting fit was good, the model was only fitted to the group averages rather than to individual participants' data.[2] In addition, the authors did not fit an exponential discounting model as a comparison. These caveats render the model fits difficult to interpret.

Investigations into dynamic inconsistency have not been entirely uniform, however. Several studies have failed to find evidence for the preference reversals expected under hyperbolic discounting. For example, Ahlbrecht and Weber (1997; see also Kable & Glimcher, 2010) manipulated front-end delay similar to the studies reviewed above. The authors then tallied the number of participants that exhibited increasing, decreasing, or equivalent patience as the front-end delay was increased. Surprisingly, the majority of participants did not shift their preferences (for similar absence of preference reversals, see Kable & Glimcher, 2010; Read, 2001). Read and Roelofsma (2003) employed a slightly different method in which participants supplied an immediate reward that would be equivalent to a temporally delayed reward (what Ahlbrecht & Weber, 1997, have referred to as a matching procedure). More violations of stationarity were observed when using matching than when using more traditional choice proce-

[2] Of course, one should always be cautious when fitting models to aggregate data because there is no guarantee that averages computed across participants will resemble any of the individual participants' data. However, in this particular case, it is even more problematic because the average of exponential curves has been shown to be a hyperbolic curve (Kurth-Nelson & Redish, 2009). Thus, the analyses of Holt et al. (2008) suggest that the group behavior resembled exactly what would be expected if each individual decision-maker were to discount exponentially.

dures. However, their data suggest that, even when participants did exhibit inconsistent preferences, the observed preferences did not match the predictions of hyperbolic discounting. Instead, the observed preferences were “actually far more constant than predicted by the conventional hyperbolic discounting model” (Read & Roelofsma, 2003, p. 148). That is, participants became more patient when considering rewards further and further in the future (a violation of stationarity), but not nearly to the degree expected under hyperbolic discounting.

To summarize, the exponential discount function (Equation 1), which is the economically normative method of discounting, has found relatively little empirical support. Instead, researchers have argued with near uniformity that decision-makers’ temporal preferences are hyperbolic (Equation 2) in nature. Consistent with this claim, individual discount functions do appear to be well fit by Equation 2 (e.g., McKerchar et al., 2009). More importantly, many (though not all) of those studies that have looked for direct evidence of preference reversals have found that preferences do shift in the direction expected under hyperbolic discounting; decision-makers tend to become more impatient as delayed rewards draw near.

Given the strong normative standard provided by economics and the apparent deviations exhibited by actual decision-makers, the large majority of previous empirical work has focused on refuting exponential discounting as a descriptive account of behavior. In addition, because of the strong theoretical grounding of hyperbolic discounting (e.g., Chung & Herrnstein, 1967), evidence against exponential discounting has largely been accepted as prima facie evidence for hyperbolic discounting. It is these factors, presumably, that have contributed to previous studies’ relatively lax evaluation of the specific, quantitative predictions of hyperbolic discounting. Unfortunately, this focus has led to a gap in our knowledge. That is, observing violations of stationarity falsifies exponential discounting as a valid empirical description, but does not provide particularly good evidence in favor of hyperbolic discounting.

The current study is designed to more directly assess delay discounting behavior, specifically focusing on the preference reversals predicted by hyperbolic discounting. However, unlike previous studies, I am not simply interested in whether or not preference reversals occur. Instead, I am interested in whether preference reversals are consistent with the quantitative predictions of hyperbolic discounting. To do so, I employed a traditional front-end delay paradigm with one important change. As the front-end delay was manipulated, the larger of the two rewards was also manipulated. Specifically, these rewards were altered so as to precisely counteract the influence of the front-end delay manipulation for a hyperbolic discounter. For such a decision-maker, increases in the front-end delay will tend to make the larger, more delayed reward more attractive, so decreasing the larger of the two rewards makes it possible to nullify this tendency. In this way, I was able to construct sets of items that should elicit uniform choices from a decision-maker discounting according to Equation 2. In contrast, any systematic shifts in choice behavior would indicate preferences that were inconsistent with hyperbolic discounting. This paradigm will hold hyperbolic discounting to the same empirical standard that has been applied to exponential discounting.

Method

Participants

Participants were 51 Stony Brook University undergraduate psychology students participating for partial course credit.

Design and Procedure

The task consisted of 108 questions, each of which asked participants to choose between sums of money available at some specified time. For example, one of the items was a choice between $40 delivered tonight and $55 delivered in 25 days. Twenty-one of these items have previously been used (Kirby & Marakovic, 1996) to quantify decision-makers’ discounting rates. These items always consist of an immediate reward (available “tonight”) and a delayed reward. The size of the monetary rewards and the delay vary across these 21 items so as to allow estimates of a wide variety of discount rates. Assuming hyperbolic discounting, these items are sufficient to detect discount rates (k in Equation 1) from 0.0007 to 0.1333. In my experience, I have found that undergraduates’ preferences fall toward the impatient end of this range. To ensure that I did not artificially exclude particularly impatient participants, I amended the original 21 items with 6 additional items that extended the range of measurable discounting rates from 0.0007 to 1.0.

The remaining 81 items were created by adding a front-end delay of either 10, 20, or 30 days to each of the original 27 items. The four sets of items allowed me to re-estimate participants’ discount rates for each of the four different front-end delays. Importantly, the four sets of 27 items were additionally modified so as to each elicit an identical pattern of choices according to hyperbolic discounting (Equation 2). Because hyperbolic discounting is sensitive to both the relative and the absolute delays of the reward pairs, I manipulated both delay and reward in order to equate the predicted pattern of choices across the four sets of items. Specifically, the larger reward was set to be

$$V_{SS}\left(\frac{1 + k_{indiff}D_{LL}}{1 + k_{indiff}D_{SS}}\right) \tag{3}$$

where V_{SS} is the smaller reward, k_{indiff} is the discount rate that would make a hyperbolic discounter indifferent between the delayed and immediate rewards associated with that item when the front-end delay was zero and D_{LL} and D_{SS} were the delays associated with the larger and smaller rewards associated with that item under the new, modified front-end delay. This new, modified reward was then rounded to the nearest whole dollar. For example, participants would be asked to choose between $40 delivered tonight and $55 delivered in 25 days, but also asked to choose between $40 delivered in 10 days and $53 delivered in 35 days, between $40 delivered in 20 days and $52 delivered in 45 days, and between $40 delivered in 30 days and $50 delivered in 55 days. As can be seen from this example, as the reward pair was moved further into the future, the larger reward was reduced. This reduction is such that a hyperbolic discounter that was indifferent between the first pair of rewards ($k = 0.015$ in this particular example) would also be indifferent between the other, more delayed versions of this item. It follows that a hyperbolic discounter preferring the larger of the original two rewards (implying that $k < 0.015$) would also prefer the larger reward in the other, more delayed versions of this item. Likewise, a hyperbolic discounter preferring the smaller of the original two rewards (implying that $k > 0.015$) would also prefer the smaller reward in the other, more delayed versions of this item. Thus, this is a paradigm that can be used to evaluate the influence of front-end delay on preferences expected under hyperbolic discounting. Consequently, any systematic shifts in preference seen across the four sets of items would imply that decision-makers are not expressing preferences consistent with hyperbolic discounting.

The 108 items were presented in a random order. The question, "Which would you prefer?" appeared at the top of the computer screen. Below this, the two rewards were presented on the left and right halves of the computer screen (rewards were randomly assigned to left/right). Participants used left and right arrow keys to indicate their choice.

Results

Using the procedure described by Kirby and Marakovic (1996), I estimated hyperbolic discount rates separately for the four sets of items (i.e., 0, 10, 20, and 30 day front-end delays).[3] The estimated discount rates are shown in Figure 1. These estimated discount rates were submitted to a one-way, repeated measure analysis of variance (ANOVA). The main effect of front-end delay was highly significant, $F(1, 50) = 50.13$, $p < .0001$. Indeed, the estimated discount rates for the four sets of items were all found to differ from each other, $ts(50) > 2.42$, $ps < .05$, such that longer front-end delays elicited larger discount rates. I also wished to evaluate choices in a somewhat more theoretically agnostic manner. To do so, I simply computed the proportion of trials on which participants selected the larger, later reward. A one-way, repeated measure ANOVA again revealed a main effect of front-end delay, $F(1, 50) = 53.59$, $p < .0001$. Indeed, the tendency to select the larger reward within the four sets of items were all found to differ from each other, $ts(50) > 3.00$, $ps < .005$, such that longer front-end delays were associated with fewer choices for the large reward. This pattern suggests that, as front-end delay was increased, participants did not become as patient as would be expected under hyperbolic discounting.

General Discussion

The current study sought to investigate the nature of decision-makers' discounting behavior. Specifically, I attempted to evaluate decision-makers' tendency to exhibit patient preferences when considering rewards in the distant future, but exhibit impatient preferences when considering more immediate rewards. These inconsistent preferences are expected, at least in principle, if decision-makers discount delayed rewards hyperbolically, but not if they discount exponentially. Surprisingly little work has been dedicated to these inconsistencies themselves, despite the fact that the resulting shifts in preferences are the very reason that economists find hyperbolic discounting so unsatisfactory. What work has been reported (e.g., Green et al., 1994; Holt et al., 2008; Kirby & Herrnstein, 1995) has largely found that decision-makers do indeed exhibit preference reversals in violation of the stationarity axiom. These findings have then led researchers to assume, often implicitly, that the observed behavior must therefore be consistent with hyperbolic discounting. However, preference reversals do not imply that decision-makers are discounting rewards hyperbolically, they simply indicate that decision-makers are not discounting exponentially.

The current study sought to hold hyperbolic discounting to the same empirical standard that exponential discounting has been held to. That is, I employed a standard delay discounting task and manipulated front-end delay in such a way that the choices of a hyperbolic discounter should have been unaffected (cf. Rachlin & Green, 1972). Given that the participants' choices varied systematically with front-end delay, I conclude that they must not discount hyperbolically. I do not, however, necessarily conclude that the participants discounted exponentially. Instead, it seems more plausible that individuals' preferences do exhibit dynamic inconsistency (violating the axiom of stationarity), but perhaps not as substantially as expected under hyperbolic discounting (Read & Roelofsma, 2003). Alternatively, it is possible that dynamic inconsistency is neither as pervasive as would be expected under hyperbolic discounting and/or that dynamic inconsistency reflects mechanisms that are unrelated to hyperbolic discounting (Read, Frederick, & Airoldi, 2012).

The doubt cast on hyperbolic discounting is particularly important because previous demonstrations of non-exponential discounting have spurred an enormous amount of influential work in economics (e.g., Laibson, 1997; Read, 2001; Rubinstein, 2003) including numerous theoretical attempts to rationalize hyperbolic discounting (Equation 2) specifically. For example, Sozou (1998) assumed that preferences regarding delayed rewards are based on the risk implied by the associated delay. The discount function is then a consequence of the risk assumed by the decision-maker. Sozou demonstrated that if risk decreases over time, then decision-makers will (and should) discount hyperbolically (see also Azfar, 1999). Alternatively, Kurth-Nelson and Redish (2009) have argued that hyperbolic discounting may arise because decision-makers are uncertain of their own discount rate. They demonstrate that if decision-makers' preferences are actually based on multiple exponential discount functions, each with a different discount rate, then the decision-maker will exhibit precisely hyperbolic discounting. Others (Ray & Bossaerts, 2011; Takahashi, Oono, & Radford, 2008) have argued that the non-linear, psychological representation of time may explain non-exponential discounting. For example, Takahashi et al. (2008) have demonstrated that an exponential discounter that represents time logarithmically would appear hyperbolic. Thus, these authors have argued that the diminishing patience exhibited by decision-makers may be entirely due to a non-linear representation of time. All of these proposals (and countless others from economics) represent extremely rigorous and parsimonious theoretical work. However, these proposals obviously lose some of their appeal if hyperbolic discounting is not an accurate description of decision-makers' preferences.

[3] This procedure consists of looking for "crossover" points along the continuum of discount rates represented by each set of items. I have also employed standard econometric model fitting procedures to estimate participants' discount rates. These more sophisticated methods yield qualitatively identical results.

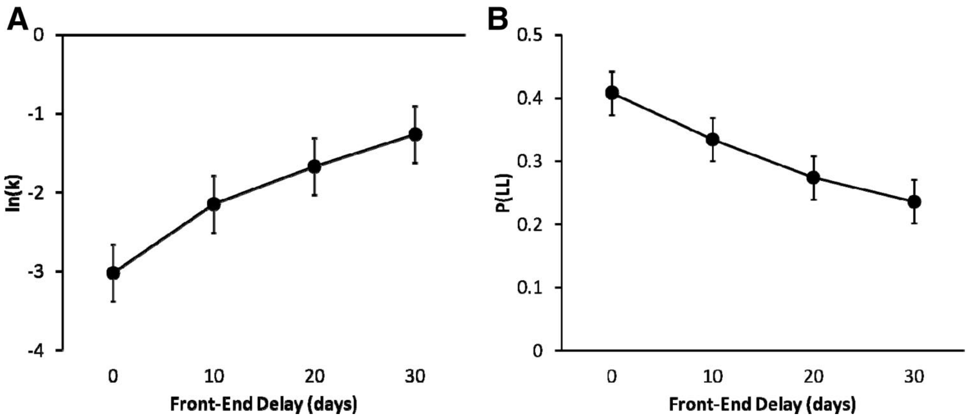

Figure 1. (A) Log-transformed hyperbolic discount rates (k) estimated from individual participants' data. (B) Probability of selecting the larger, later reward (LL). All error bars illustrate within-subject confidence intervals as suggested by Loftus and Masson (1994).

Alternative Discounting Schemes

Throughout this article, I have focused exclusively on the exponential and hyperbolic discounting functions. This is due to the literature's overwhelming emphasis on these two functional forms. However, there has been discussion of alternative discounting schemes. For example, a class of discounting functions referred to as hyperboloids (Myerson & Green, 1995; Rachlin, 2006; Rodriguez & Logue, 1988) represent modifications of traditional hyperbolic discounting (Equation 2). These proposals suggest raising the delay parameter (Rodriguez & Logue, 1988) or the entire denominator of Equation 2 (Myerson & Green, 1995) to an exponent smaller than one in order to reflect diminishing sensitivity to increasing delay. Each of the hyperboloids, with their exponents set to less than one (the typical value when fitted to empirical data; McKerchar et al., 2009), suggests that participants' should exhibit even greater violations of stationarity than implied by standard hyperbolic discounting. That is, hyperboloid discounters should have exhibited increasingly patient behavior as the front-end delay was increased in the current study. The participants exhibited an effect in exactly the opposite direction, suggesting that hyperboloid functions are not a satisfactory explanation of the current results.

Alternatively, Read and colleagues (Read, 2001; Read & Roelofsma, 2003) have proposed that violations of stationarity may be explained by appealing to an effect called subadditivity rather than hyperbolic discounting. Read and Roelofsma (2003) noted that existing empirical data has difficulty distinguishing between these qualitatively different explanations because many past studies have confounded front-end delay with the difference between the two reward delivery times. Given that the current study manipulated front-end delay independently, my front-end delay manipulations should not have had any effect according to the subadditivity explanation. More recently, Scholten and Read (2010) have proposed that intertemporal choice anomalies may arise from a similarity-based evaluation of delayed rewards rather than discount functions themselves (for related suggestions, see Leland, 2002; Rubinstein, 2003). These authors suggested that increasing front-end delay increases patience because it increases the similarity between delays (e.g., the difference between today and 12 months from now seems greater than the difference between 12 months from now and 24 months from now), not because of decision-makers' discount functions. Under this view, the current results place strong constraints on the similarity function used to evaluated delayed rewards.

References

Ahlbrecht, M., & Weber, M. (1997). An empirical study on intertemporal decision making under risk. *Management Science, 43,* 813–826. doi:10.1287/mnsc.43.6.813

Ainslie, G. (2001). *Breakdown of will.* doi:10.1017/CBO9781139164191

Ainslie, G., & Herrnstein, R. J. (1981). Preference reversal and delayed reinforcement. *Animal Learning & Behavior, 9,* 476–482. doi:10.3758/BF03209777

Azfar, O. (1999). Rationalizing hyperbolic discounting. *Journal of Economic Behavior & Organization, 38,* 245–252. doi:10.1016/S0167-2681(99)00009-8

Chung, S. H., & Herrnstein, R. J. (1967). Choice and delay of reinforcement. *Journal of the Experimental Analysis of Behavior, 10,* 67–74. doi:10.1901/jeab.1967.10-67

Green, L., & Estle, S. J. (2003). Preference reversals with food and water reinforcers in rats. *Journal of the Experimental Analysis of Behavior, 79,* 233–242. doi:10.1901/jeab.2003.79-233

Green, L., Fisher, E. B., Perlow, S., & Sherman, L. (1981). Preference reversal and self control: Choice as a function of reward amount and delay. *Behaviour Analysis Letters, 1,* 43–51.

Green, L., Fristoe, N., & Myerson, J. (1994). Temporal discounting and preference reversals in choice between delayed outcomes. *Psychonomic Bulletin & Review, 1,* 383–389. doi:10.3758/BF03213979

Holt, D. D., Green, L., Myerson, J., & Estle, S. J. (2008). Preference reversals with losses. *Psychonomic Bulletin & Review, 15,* 89–95. doi:10.3758/PBR.15.1.89

Kable, J. W., & Glimcher, P. W. (2010). An "as soon as possible" effect in human inter-temporal decision making: Behavioral evidence and neural mechanisms. *Journal of Neurophysiology, 103,* 2513–2531. doi:10.1152/jn.00177.2009

Kirby, K. N. (1997). Bidding on the future: Evidence against normative discounting of delayed rewards. *Journal of Experimental Psychology: General, 126,* 54–70. doi:10.1037/0096-3445.126.1.54

Kirby, K. N., & Herrnstein, R. J. (1995). Preference reversals due to myopic discounting of delayed reward. *Psychological Science, 6,* 83–89. doi:10.1111/j.1467-9280.1995.tb00311.x

Kirby, K. N., & Marakovic, N. N. (1996). Delay-discounting probabilistic rewards: Rates decrease as amounts increase. *Psychonomic Bulletin & Review, 3,* 100–104. doi:10.3758/BF03210748

Koopmans, T. C. (1960). Stationary ordinal utility and impatience. *Econometrica, 28,* 287–309. doi:10.2307/1907722

Kurth-Nelson, Z., & Redish, A. D. (2009). Temporal-difference reinforcement learning with distributed representations. *PLoS ONE, 4*(10), e7362. doi:10.1371/journal.pone.0007362

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

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

Loftus, G. R., & Masson, M. E. J. (1994). Using confidence intervals in within-subject designs. *Psychonomic Bulletin & Review, 1,* 476–490. doi:10.3758/BF03210951

McKerchar, T. L., Green, L., Myerson, J., Pickford, T. S., Hill, J. C., & Stout, S. C. (2009). A comparison of four models of delay discounting in humans. *Behavioural Processes, 81,* 256–259. doi:10.1016/j.beproc.2008.12.017

Millar, A., & Navarick, D. J. (1984). Self-control and choice in humans: Effects of video game playing as a positive reinforcer. *Learning and Motivation, 15,* 203–218. doi:10.1016/0023-9690(84)90030-4

Myerson, J., & Green, L. (1995). Discounting of delayed rewards: Models of individual choice. *Journal of the Experimental Analysis of Behavior, 64,* 263–276. doi:10.1901/jeab.1995.64-263

Rachlin, H. (2006). Notes on discounting. *Journal of the Experimental Analysis of Behavior, 85,* 425–435. doi:10.1901/jeab.2006.85-05

Rachlin, H., & Green, L. (1972). Commitment, choice and self-control. *Journal of the Experimental Analysis of Behavior, 17,* 15–22. doi:10.1901/jeab.1972.17-15

Ray, D., & Bossaerts, P. (2011). Positive temporal dependence of the biological clock implies hyperbolic discounting. *Frontiers in Decision Neuroscience, 5,* 2. doi:10.3389/fnins.2011.00002

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

Read, D., Frederick, S., & Airoldi, M. (2012). Four days later in Cincinnati: Longitudinal tests of hyperbolic discounting. *Acta Psychologica, 140,* 177–185. doi:10.1016/j.actpsy.2012.02.010

Read, D., & Roelofsma, P. H. M. P. (2003). Subadditive versus hyperbolic discounting: A comparison of choice and matching. *Organizational Behavior and Human Decision Processes, 91,* 140–153. doi:10.1016/S0749-5978(03)00060-8

Rodriguez, M. L., & Logue, A. W. (1988). Adjusting delay to reinforcement: Comparing choice in pigeons and humans. *Journal of Experimental Psychology: Animal Behavior Processes, 14,* 105–117. doi:10.1037/0097-7403.14.1.105

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

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

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

Solnick, J. V., Kannenberg, C. H., Eckerman, D. A., & Waller, M. B. (1980). An experimental analysis of impulsivity and impulse control in humans. *Learning and Motivation, 11,* 61–77. doi:10.1016/0023-9690(80)90021-1

Soman, D., Ainslie, G., Frederick, S., Li, X., Lynch, J., Moreau, P., . . . Zauberman, G. (2005). The psychology of intertemporal discounting: Why are distant events valued differently from proximal ones? *Marketing Letters, 16,* 347–360. doi:10.1007/s11002-005-5897-x

Sozou, P. D. (1998). On hyperbolic discounting and uncertain hazard rates. *Proceedings of the Royal Society of London B, 265,* 2015–2020. doi:10.1098/rspb.1998.0534

Takahashi, T., Oono, H., & Radford, M. H. B. (2008). Psychophysics of time perception and intertemporal choice models. *Physica A, 387,* 2066–2074. doi:10.1016/j.physa.2007.11.047

Received August 16, 2012
Revision received September 24, 2012
Accepted September 25, 2012 ■