

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Subjective Hazard Rates Rationalize “Irrational” Temporal Preferences

Permalink

<https://escholarship.org/uc/item/8q0511km>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 36(36)

ISSN

1069-7977

Authors

Luhmann, Christian
Bixter, Michael

Publication Date

2014

Peer reviewed

Subjective Hazard Rates Rationalize “Irrational” Temporal Preferences

Christian C. Luhmann (christian.luhmann@stonybrook.edu)

Department of Psychology
Stony Brook, NY 11794 USA

Michael T. Bixter (michael.bixter@stonybrook.edu)

Department of Psychology
Stony Brook, NY 11794 USA

Abstract

Delay discounting refers to decision makers’ tendency to treat immediately consumable goods as more valuable than those only available after some delay. Previous work has focused on a seemingly irrational feature of these preferences: the systematic tendency to exhibit more patience when consequences are far in the future but less patience about those same, identical rewards as time passes. One explanation for delay discounting itself appeals to the risk implicitly associated with delayed rewards. The current study investigates whether the implicit risk hypothesis is capable of explaining the seemingly irrational shifts in patience by having participants make subjective risk judgments regarding a variety of real-world scenarios. To reduce the possibility of task demands, participants judged hazard rates rather than survival rates. Results suggest that the seemingly irrational shifts in patience are quite reasonable once participants’ beliefs about the relationship between delay and risk are taken into account.

Keywords: discounting; implicit risk; hazard function

People must frequently make choices that involve temporal considerations. For example, would you prefer a job with the higher starting salary or a job with greater opportunity for future advancement? Prior work on such temporal decisions has primarily focused on exploring the phenomenon of delay discounting, in which decision makers behave as though immediately consumable goods are more valuable than those only available after some delay. This phenomenon is powerful enough that decision makers are frequently willing to forgo delayed rewards in favor of immediate rewards even when the delayed rewards are objectively more valuable. For example, a decision maker might choose an immediately-available \$100 over \$200 that would only be delivered in three years. Such a preference is assumed to reflect the subjective value of the \$200 option, discounted by the associated three-year delay. The sway of negative events is similarly blunted by delay. For example, working on your taxes next month is likely preferred over working on them tonight.

Classical economics has viewed delay discounting from within the framework of discounted utility theory (Samuelson, 1937), according to which the subjective value of goods drops exponentially:

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

where V_D represents the current value of a reward that will be delivered at time D , k represents the decision maker’s

discount rate, and V_0 represents the undiscounted subjective value of that same reward (i.e., the value of that reward if it were available immediately). This expression suggests that the subjective value drops by a fixed percentage for each unit of time that those goods are delayed. If a decision maker’s discount rate is 0.1 (and D is measured in years), then \$100 available in a year is only worth \$90 right now, or 90% of its immediate value. This reward delayed an additional year is only worth \$81, reflecting an additional 10% drop in value.

Given that exponential discounting represents an obvious normative standard, there has been a large amount of work designed to carefully evaluate how the behavior of human decision makers compares (Ainslie, 1992; Ainslie & Herrnstein, 1981; Green, Fristoe, & Myerson, 1994; Kirby, 1997; Kirby & Herrnstein, 1995; Loewenstein & Thaler, 1989; Rachlin, 1995). Nearly all of this work has demonstrated that decision makers do not discount exponentially. Instead, researchers have largely advocated hyperbolic discount functions as a superior descriptive account. Under hyperbolic discounting, the subjective value of a delayed reward can be expressed as below, with the quantities defined as above.

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

The critical difficulty with hyperbolic discounting is that it tends to result in contradictory preferences. For example, a hyperbolic discounter might prefer \$200 delivered in nine years rather than \$100 delivered in six years (a patient preference) but also prefer an immediate \$100 rather than \$200 delivered in 3 years (an impatient preference). The problem, of course, is that the first pair of rewards will become the second pair of rewards in six years time. Such contradictory preferences violate the axiom of stationarity (Koopmans, 1960), and their predictable nature means that such decision makers may be exploited. That is, an exponential discounter can continually offer the hyperbolic discounter a choice between delayed pairs of rewards, only to later offer (sell) the hyperbolic discounter an opportunity to change their selection and profit from the reliable reversal of preference. For this reason, hyperbolic discounting has been characterized as irrational (Kirby, 1997; Soman et al., 2005). Real-world examples of these preference reversals are not difficult to find. For example, the preferences reflected in our New Year’s resolutions (e.g., exercising over the next 12 months) reliably shift once the time comes

to face consequences (e.g., getting up this morning to go to the gym). These preference reversals have also been blamed for a variety of undesirable patterns of behavior such as drug addiction and procrastination (Ainslie, 2001).

Implicit Risk Hypothesis

Despite the strong interest in the general phenomenon, there has been surprisingly little work exploring why decision makers exhibit delay discounting in the first place. That is, why should a decision maker behave as though \$100 were worth any less than \$100 regardless of when it would be delivered? The predominant explanation in economics (Azfar, 1999; Benzion, Rapoport, & Yagil, 1989; Dasgupta & Maskin, 2005; Prelec & Loewenstein, 1991; Sozou, 1998; Yaari, 1965), ecology (Kacelnik, 2003), and psychology (Mazur, 1989, 1995, 1997; Mischel, 1966; Rachlin, Raineri, & Cross, 1991; Stevenson, 1986) is that the delay associated with postponed consequences renders them inherently uncertain. This has been referred to as the implicit risk hypothesis (Benzion et al., 1989) and justifies the general phenomenon of delay discounting. For example, a decision maker confronted with a typical intertemporal dilemma might prefer a more immediate reward if it were believed that the delayed reward's future availability could not be guaranteed (i.e., the promisor could go bankrupt, the decision maker could die, etc.). Furthermore, if decision makers believe that longer delay intervals imply lower probabilities of receipt, then discounted value should decrease monotonically as delay increases. This hypothesis suggests that delay discounting behavior always involves considerations about risk, even when information about information about risk is omitted (as in standard intertemporal choice tasks) and even if experimenters attempt to "guarantee" rewards (e.g., it's not clear how to "guarantee" that the decision maker will survive until receipt).

Under the implicit risk hypothesis, decision makers' discounting preferences stem directly from their beliefs about the relationship between risk and delay. In particular, discounting functions can be thought of as reflecting the survival rate, $s(t)$, of the delayed reward. The survival rate specifies the probability with which a promised reward "survives" until time t . Alternatively, it can be thought of as the probability with which a promised reward can be successfully consumed at time t . Intuitively, higher survival rates should encourage smaller discount rates (patience) whereas lower survival rates should encourage larger discount rates (impatience). Frequently, it is more convenient to instead work with the hazard function,

$$h(t) = -\frac{1}{s} \frac{ds}{dt} \quad (3)$$

which specifies the probability of a reward being lost between time t and time $t+dt$, given that it has survived until time t . The hazard function is particularly useful when assessing the shape of the discount function (Sozou, 1998). For example, under the implicit risk hypothesis, exponential discounting suggests a constant hazard rate, $h(t)=\lambda$. That is, exponential discounting suggests that decision makers'

beliefs about risk are independent of delay (see Figure 1). This explains why exponential discounters exhibit stationary preferences. The probability of a delayed reward disappearing between now and 12 months from now is identical to the probability of a delayed reward disappearing between 12 months from now and 24 months from now. As a consequence, if \$200 in two years is preferred to \$100 in a year, \$200 in a year will also be preferred over \$100 immediately.

In contrast, hyperbolic discounting suggests at least two different conceptions of risk. First, hyperbolic discounting is consistent with the belief that the hazard rate declines over time (see Figure 1). That is, hyperbolic discounting may suggest that the probability of a delayed reward disappearing between now and 12 months from now is greater than the probability of a delayed reward disappearing between 12 months from now and 24 months from now. The second possibility is that hyperbolic discounters may believe a constant hazard rate, $h(t)=\lambda$, but may be uncertain about the true value of λ (Azfar, 1999). For example, imagine a decision maker offered a choice between an immediate \$100 and \$200 to be delivered in a year. A decision maker might be uncertain about the hazard rate (i.e., it might be quite high) and thus might prefer the immediate \$100. Now imagine that this same decision maker is faced with a choice between \$100 in a year and \$200 in two years. Should the decision maker still prefer the more immediate \$100? It is conceivable that, if the true hazard rate is low enough for the \$100 to survive for 12 months, it might also be low enough for the \$200 to survive 24 months. Such reasoning would thus justify the violations of stationarity underlying many of decision makers' inconsistent preferences (though not all, Dasgupta & Maskin, 2005).

Perceived Risk

Non-exponential discounting has been deemed irrational because of the contradictory preferences it generates. However, the dominant explanation for discounting, the implicit risk hypothesis, suggests that discounting should strongly depend on decision makers' beliefs about the risk they face. If decision makers believe in declining hazard functions, or if they are uncertain about the true hazard rate, then non-exponential discounting may actually be justifiable. Thus, direct empirical evaluation of decision makers' beliefs about risk could illuminate the extent to which violations of stationarity, and hyperbolic discounting more broadly, are rational.

Recent work has begun to investigate the quantitative relationship between perceived risk and delay discounting. For example, Patak and Reynolds (2007) asked participants to assess the certainty associated with various delayed rewards. Results indicated that perceived certainty declined as the length of the delay intervals increased. In addition, participants' certainty judgments were predictive of their discount rates. Takahashi, Ikeda, and Hasegawa (2007) used a similar method to investigate whether subjective risk

could explain non-exponential discounting. Not only did results indicate that more delayed rewards were less certain, but the subjective probability of delayed rewards was found to drop hyperbolically as a function of delay. Thus, the subjective probability results were exactly what the implicit risk hypothesis would predict.

There are several reasons to treat these findings with caution, however. For example, Takahashi et al. (2007) employed a within-subject design in which participants always completed the subjective probability questionnaire after completing the choice task and were explicitly instructed to consider their earlier value judgments when making their probability judgments. Thus, the strong relationship between the subjective value and discounted value may have been induced by the task itself.

The current study is designed to provide an empirical evaluation of the relationship between delay discounting and the beliefs about risk in the real world. However, given our concerns about prior, similar investigations, the current study takes a different strategy. Instead of assessing the survival rate, participants in the current study provide judgments of the hazard rate; the probability of a reward being lost between time t and time $t+dt$ given that it has already survived an interval of length t . If participants believe that the hazard rate is constant, judgments should be insensitive to the specific value of t . If, on the other hand, decision makers believe that the hazard rate drops over time, then we would expect judgments to increase as t increases. Similarly, if decision makers are uncertain about the true value of the underlying hazard rate, then we would also expect probability judgments to increase as t increases. This latter prediction is due to the idea that the continued survival of delayed rewards provides information about the hazard rate, suggesting that it may be lower than previously believed (Azfar, 1999; Sozou, 1998). Finally, we believe that evaluating hazard rates rather than survival rates reduces the possibility of task demands because the nature of the relationship between a given discounting function and its corresponding hazard function is unlikely to be obvious to participants (whereas discounting functions and survival functions are essentially equivalent).

Method

Participants

Forty-three Stony Brook University undergraduates participated for partial course credit. One participant exclusively used the extremes of the response scale (all probabilities were judged to be either 0 or 1) and was not included in subsequent analyses.

Probability Judgment Task

The study consisted of a set of 36 probability judgments, each concerning six scenarios. Each judgment asked participants about the probability of some proposition being true at some specific time in the future. For each scenario, we constructed three different conditional probability

judgments. Each conditional probability judgment requested a probability of the form, “How likely will X continue to be true for the next $t + dt$ given that it has been true for the last t ?” where X was some proposition and t and dt were temporal intervals (e.g., 3 weeks). This judgment provides information about the hazard rate associated with the interval from t to $t+dt$. Specifically, this judgment gives us the complement of the hazard rate (i.e., $1 - h_{t+dt}$). The three conditional probabilities differed only in the value of t . Each scenario had a Short, Medium, and Long version, each of which employed a large, medium, or small value t . Importantly, the value of dt used for each scenario was such that $dt = t_{long} - t_{med} = t_{med} - t_{short}$.

Instead of asking for the conditional probability judgments in isolation, these judgments were always immediately preceded by a related, unconditional probability judgment. For example, one item asked 1) how likely it was that an open seat in a college course would still be open in one week and 2) how likely it was that an open seat would still be open knowing that it had been open for the last 12 weeks¹. The first, unconditional probability judgment represented an assessment of the survival rate. That is, this judgment requested the probability that the proposition was true at time t given no information about how long that proposition had been true for (alternatively, one can interpret this as the probability that the proposition would continue to be true for the next $t + dt$ given that it has been true for the last t , where t is interpreted to be zero). Importantly, the unconditional probability judgment always queried the survival function at time $t = t_{short}$. This, along with the specific values of t_{short} , t_{med} , t_{long} , and dt allow us to reconstruct the full survival function (described below). Including the unconditional judgments also ensured that participants did not misinterpret the request for conditional probability judgments as requests for unconditional probability judgments. To emphasize the fact that each unconditional and conditional judgment formed a pair, the inter-trial interval (ITI) between the two halves of the pair was 500ms whereas the ITI between two different pairs was 4 seconds.

The items were sequenced such that each of the six scenarios was presented once in each of three blocks. The order in which the different scenarios had their Long, Medium, and Short conditional probability judgments made was pseudorandom and counterbalanced across participants. In the first block, each participant would see each of the six scenarios, making two Long conditional probability judgments, two Medium conditional probability judgments, and two Short conditional probability judgments. The remaining blocks were similar, allowing each participant to provide all judgments for all scenarios. Because the unconditional probability judgment always preceded the conditional probability judgment, this meant that each

¹ Other scenarios concerned: whether a specific movie was showing, whether an employment opportunity was filled, whether a coupon was valid, whether tickets to the zoo were available, and whether an item was on sale at Macy’s.

participant made the unconditional probability judgment for each scenario three times (once per block). For the purposes of data analysis, we only considered the first unconditional judgment from each scenario. This way, the unconditional judgments are uncontaminated by any subsequent judgments.

Discounting Task

The discounting task consisted of 12 trials, each of which presented two rewards. Each reward included a magnitude (in dollars) and a delay until the reward would be received (in months). Importantly, each trial omitted the reward magnitude of the sooner option. Participants' task was to supply this missing reward magnitude with a value that would render them indifferent between the two options. Half the trials included an immediate reward (e.g., what amount delivered today would be equivalent to \$30 delivered in 6 months?) and half involved two delayed rewards (e.g., what amount delivered in 3 months would be equivalent to \$75 delivered in 9 months?).

Procedure

After completing consent paperwork and receiving brief instructions, participants completed a single practice item. The item allowed participants to see how each probability judgment would be paired with a related, conditional probability judgment. It also allowed practice with the probability response scale which was a visual analog scale with left end labeled as, "Very Unlikely" and the right end labeled as, "Very Likely", but otherwise unmarked. Participants used the left and right arrow keys to move a cursor to some location on the scale and pressed the enter key to move on to the next item. Responses were converted to probabilities by scaling the final cursor position such that zero represented the left end of the scale (very unlikely) and one represented the right end of the scale (very likely). The entire procedure took approximately 15 minutes.

Results

To explore participants' judgments (Figure 1), we performed a six (scenario) by four (judgment: Unconditional, Short, Medium, Long) repeated measures ANOVA. As predicted, we observed a main effect of judgment type ($F(3, 126)=6.96, p<.0001$). We also observed a significant interaction between scenario and judgment ($F(3, 126)=1.34, p<.0001$). The main effect of judgment was due to the fact that unconditional judgments ($M=.44, SD=.12$) were less than conditional judgments ($M=.50, SD=.12, t(41)=2.48, p<.05$). Furthermore, the conditional judgments in the Long condition ($M=.53, SD=.16$) were higher than those in the Medium condition ($M=.50, SD=.13, t(41)=2.79, p<.01$) and those in the Short condition ($M=.48, SD=.12, t(41)=2.17, p<.05$). Judgments in the Medium were greater than those in the Short condition, though not significantly so.

The fact that these probability judgments increased as t increased is consistent with a decreasing hazard rate; the

longer the reward has survived, the more likely it was expected to survive an additional dt delay. This pattern is also consistent with Sozou's (1998) suggestion that prolonged survival allows decision makers to revise (i.e., lower) their uncertain estimates of a fixed hazard rate. In either case, participants' judgments are entirely consistent with the violations of stationarity typically observed in discounting behavior.

Characterizing Survival and Discount Functions

To further illustrate the implications for delay discounting, we reconstructed the survival function (i.e., discount functions) implied by participants' probability judgments. We first assumed that $s(0) = 1.0$. As described in the method section, the unconditional judgments provided $s(t_{short})$. The conditional probability judgments provided information about the hazard function, h , which allowed us to estimate three additional points of the survival function. For example, $s(t_{short} + dt)$ can be estimated as $s(t_{short}) \cdot [1 - h_{short+dt}]$, where $[1 - h_{short+dt}]$ is provided by the Short conditional probability judgment. In other words, the probability of the reward surviving the interval $[0, t_{short} + dt]$ requires the reward to first survive the initial period, t_{short} , (which occurs with probability $s(t_{short})$) and additionally survive the next interval, dt (which is judged to occur with probability $[1 - h_{short+dt}]$). Similarly, $s(dt + t_{med})$ can be estimated as $s(dt + t_{short}) \cdot [1 - h_{t_{med}+dt}]$ and $s(t_{long} + dt)$ can be estimated as $s(t_{med} + dt) \cdot [1 - h_{t_{long}+dt}]$. These quantities were computed separately for each scenario and for each participant.

Thus, each participant had six survival functions, reconstructed as described above, plus their judgments from the delay discounting task. Each of these seven sets of data was separately fitted with a generalized hyperbolic discount function.

$$V_D = V_0 \cdot \left[(1 + \theta k D)^{-\left(\frac{1}{\theta}\right)} \right] \quad (4)$$

This parametric function represents a family of discounting functions that includes both traditional exponential and hyperbolic functions as special cases. Importantly, we




Figure 1 – Average probability judgments. Judgments increased from Short to Long, a pattern consistent with commonly observed violations of stationarity.

employ a specific characterization of the generalized hyperbolic (Benhabib, Bisin, & Schotter, 2010). This characterization has a parameter, θ , that controls where along the exponential-hyperbolic continuum the resulting function falls. When $\theta = 1$, Equation 4 is equivalent to Equation 2 (the hyperbolic function). As θ approaches zero, Equation 4 approaches Equation 1 (the exponential function). Thus, the estimated value of θ quantifies the degree to which a discounting function violates the axiom of stationarity. We found values of θ and k that maximized the log-likelihood of the data, separately for each scenario and for the discounting data.

Relating Survival and Discount Functions

For each participant, we computed the median value of θ across the six survival functions. As shown in Figure 2, these median values were significantly correlated with the values of θ inferred from the discounting task (Spearman rank-order, $r = 0.37, p < .05$). This relationship suggests that violations of stationarity implied by participants' discounting behavior was explained to a large extent by their beliefs about risk. Similar analyses with estimated values of the discount rate revealed no significant relationship.

We also investigated the relationship between the survival functions estimated for each scenario. Specifically, we computed all the pairwise correlations between the values of θ inferred from each of the scenarios. The average correlation between scenarios was quite strong ($r = .22$), with coefficients ranging from 0 to .47. In contrast, similar analyses found that estimates of the discount rate were not as strongly correlated across scenarios (average $r = .15$), with many coefficients being negative (ranging from -.14 to .41). This suggests that the tendency to violate stationarity (captured by θ and rationalized by a decreasing hazard function) may reflect a relatively domain-general belief about the relationship between risk and delay in one's environment. In contrast, beliefs about the overall level of risk in the environment (reflected by the discount rate) may vary across domains.

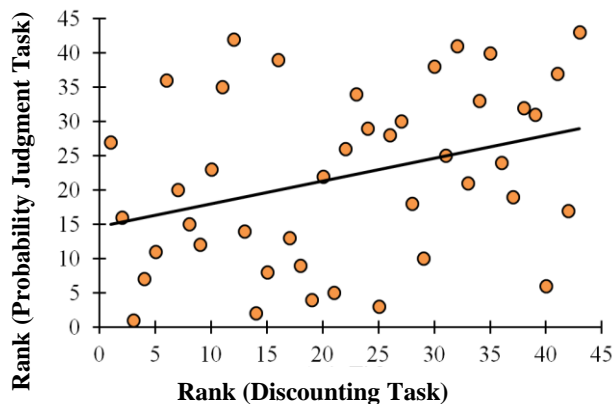


Figure 2 – Relationship between the value of theta estimated from the discounting task and the median value of theta estimated from the probability judgment task.

Discussion

In the current study we have attempted to evaluate the implicit risk hypothesis as an explanation of the non-stationary preferences typically exhibited by decision makers. We evaluated subjective beliefs about the hazard function present in real world situations. Specifically, participants were asked for the probability of a reward surviving the interval between now and time dt given that it has survived for t . Results indicate that these judgments increased as t increased. This finding is consistent with a decreasing hazard function; the longer the reward has survived, the more likely it was expected to survive an additional dt delay. Alternatively, the pattern of judgments is also consistent with Sozou's (1998) suggestion that prolonged survival would allow decision makers to revise (lower) their uncertain estimates of a fixed hazard rate.

These findings suggest a parsimonious explanation of a problematic and seemingly irrational pattern of behavior. Roughly speaking, the reason that violations of stationarity are so unpalatable to standard economic theory is because a delay of fixed length is assumed to be identical regardless of where in time this interval happens to fall. Thus, if \$200 delivered in nine years is preferred over \$100 delivered in six years (a difference of 3 years), then an immediate \$100 must also be preferred over \$200 delivered in 3 years (also a difference of 3 years). The current results suggest that there are legitimate reasons for decision makers to treat these 3-year intervals differently.

The more pragmatic reason that hyperbolic discounting has been deemed irrational is that such decision makers can be exploited by "rational", exponential discounters. However, if our participants' beliefs are accurate and hazard rates truly decline, then this arbitrage relationship actually goes the opposite direction; hyperbolic discounters can exploit exponential discounters. Alternatively, if the current results stem from our participants' uncertainty about a constant hazard rate (Azfar, 1999; Sozou, 1998), then exponential discounters will only have an advantage if they hold accurate beliefs about the risk present in their environment (i.e., they use a discount rate that is perfectly calibrated for the ambient level of risk). However it is important to note that this superiority would derive almost entirely from an informational advantage (i.e., somehow having knowledge of the true hazard rate).

The current results support the implicit risk hypothesis as a prime factor in discounting behavior and dovetail with recent work that has attempted to find direct evidence in favor of the implicit risk hypothesis. For example, Bixter and Luhmann (2014) asked decision makers in the Delay-First condition to choose between an immediate, guaranteed \$16 and a larger, delayed reward (e.g., \$44 delivered in 12 days). They were then immediately asked about the same reward pair, with the larger reward modified to now be explicitly risky (e.g., an 85% chance of receiving \$44 delivered in 12 days). In the Risk-First condition, the delay and risk information were added to the larger reward in the opposite order (e.g., the larger reward was first risky, and

then modified to be both risky and delayed). Despite the second halves of these two choices being identical choices (e.g., \$16 vs. an 85% chance of receiving \$44 delivered in 12 days), decision makers were more likely to accept the delayed, uncertain reward when the delay information was presented before the risk information than the opposite order. Not only were participants more likely to accept the rewards when presented in this order, their reaction times were faster when the delay information was presented first. Taken together, these results were interpreted to represent a sort of “priming” effect in which the contemplation of delayed rewards automatically invoked beliefs about risk, but not the other way around. These findings also reinforce a basic premise of the implicit risk hypothesis: decision makers do not believe that delayed rewards can be guaranteed, even when explicitly told that they will be delivered with 100% certainty. The current results extend this demonstration by providing information about the nature of decision makers’ uncertainty. Furthermore, because we elicited beliefs about hazard rates, rather than survival rates, we believe that the current study captured participants’ beliefs in a relatively neutral manner.

References

- Ainslie, G. (1992). *Picoeconomics: The Strategic Interaction of Successive Motivational States Within the Person*: Cambridge University Press.
- Ainslie, G. (2001). *Breakdown of Will*. Cambridge: Cambridge University Press.
- Ainslie, G., & Herrnstein, R. J. (1981). Preference reversal and delayed reinforcement. *Animal Learning and Behavior*, 9(4), 476-482.
- Azfar, O. (1999). Rationalizing hyperbolic discounting. *Journal of Economic Behavior & Organization*, 38(2), 245-252.
- Benhabib, J., Bisin, A., & Schotter, A. (2010). Present-bias, quasi-hyperbolic discounting, and fixed costs. *Games and Economic Behavior*, 69(2), 205-223.
- Benzion, U., Rapoport, A., & Yagil, J. (1989). Discount rates inferred from decisions: An experimental study. *Management Science*, 35(3), 270-284.
- Bixter, M. T., & Luhmann, C. C. (2014). Evidence for implicit risk: Delay primes the processing of uncertainty, Manuscript submitted for publication.
- Dasgupta, P., & Maskin, E. (2005). Uncertainty and hyperbolic discounting. *American Economic Review*, 95(4), 1290-1299.
- Green, L., Fristoe, N., & Myerson, J. (1994). Temporal discounting and preference reversals in choice between delayed outcomes. *Psychonomic Bulletin & Review*, 1(3), 383-389.
- Kacelnik, A. (2003). The evolution of patience. In G. Loewenstein, D. Read & R. F. Baumeister (Eds.), *Time and decision: Economic and Psychological Perspectives on Intertemporal Choice* (pp. 115-138).
- Kirby, K. N. (1997). Bidding on the future: Evidence against normative discounting of delayed rewards. *Journal of Experimental Psychology: General*, 126(1), 54-70.
- Kirby, K. N., & Herrnstein, R. J. (1995). Preference reversals due to myopic discounting of delayed reward. *Psychological Science*, 6(2), 83-89.
- Koopmans, T. C. (1960). Stationary ordinal utility and impatience. *Econometrica*, 287-309.
- Loewenstein, G., & Thaler, R. H. (1989). Anomalies: Intertemporal choice. *The Journal of Economic Perspectives*, 3(4), 181-193.
- Mazur, J. E. (1989). Theories of probabilistic reinforcement. *Journal of the Experimental Analysis of Behavior*, 51(1), 87-99.
- Mazur, J. E. (1995). Conditioned reinforcement and choice with delayed and uncertain primary reinforcers. *Journal of the Experimental Analysis of Behavior*, 63(2), 139-150.
- Mazur, J. E. (1997). Choice, delay, probability, and conditioned reinforcement. *Learning & Behavior*, 25(2), 131-147.
- Mischel, W. (1966). Theory and research on the antecedents of self-imposed delay of reward. *Progress in Experimental Personality Research*, 3, 85-132.
- Patak, M., & Reynolds, B. (2007). Question-based assessments of delay discounting: Do respondents spontaneously incorporate uncertainty into their valuations for delayed rewards? *Addictive behaviors*, 32(2), 351-357.
- Prelec, D., & Loewenstein, G. (1991). Decision Making over Time and under Uncertainty: A Common Approach. *Management Science*, 37(7), 770-786.
- Rachlin, H. (1995). Self-control: Beyond commitment. *Behavioral and Brain Sciences*, 18(1), 109-159.
- Rachlin, H., Raineri, A., & Cross, D. (1991). Subjective probability and delay. *Journal of the Experimental Analysis of Behavior*, 55(2), 233-244.
- Samuelson, P. A. (1937). A note on measurement of utility. *The Review of Economic Studies*, 155-161.
- Soman, D., Ainslie, G., Frederick, S., Li, X., Lynch, J., Moreau, P., et al. (2005). The psychology of intertemporal discounting: Why are distant events valued differently from proximal ones? *Marketing Letters*, 16(3), 347-360.
- Sozou, P. D. (1998). On hyperbolic discounting and uncertain hazard rates. *Proceedings of the Royal Society of London B*, 265(1409), 2015-2020.
- Stevenson, M. K. (1986). A discounting model for decisions with delayed positive or negative outcomes. *Journal of experimental psychology. General*, 115(2), 131-154.
- Takahashi, T., Ikeda, K., & Hasegawa, T. (2007). A hyperbolic decay of subjective probability of obtaining delayed rewards. *Behavioral and Brain Functions*, 3(1), 52.
- Yaari, M. E. (1965). Uncertain lifetime, life insurance, and the theory of the consumer. *The Review of Economic Studies*, 137-150.