

Temporal Discounting Predicts Risk Sensitivity in Rhesus Macaques

Benjamin Y. Hayden^{1,2,*} and Michael L. Platt^{1,2,3,4}

¹Department of Neurobiology

²Center for Neuroeconomic Studies

Duke University Medical School

Durham, North Carolina 27710

³Department of Biological Anthropology and Anatomy

⁴Center for Cognitive Neuroscience

Duke University

Durham, North Carolina 27710

Summary

Humans and animals tend both to avoid uncertainty and to prefer immediate over future rewards. The comorbidity of psychiatric disorders such as impulsivity, problem gambling, and addiction suggests that a common mechanism may underlie risk sensitivity and temporal discounting [1–6]. Nonetheless, the precise relationship between these two traits remains largely unknown [3, 5]. To examine whether risk sensitivity and temporal discounting reflect a common process, we recorded choices made by two rhesus macaques in a visual gambling task [7] while we varied the delay between trials. We found that preference for the risky option declined with increasing delay between sequential choices in the task, even when all other task parameters were held constant. These results were quantitatively predicted by a model that assumed that the subjective expected utility of the risky option is evaluated based on the expected time of the larger payoff [5, 6]. The importance of the larger payoff in this model suggests that the salience of larger payoffs played a critical role in determining the value of risky options. These data suggest that risk sensitivity may be a product of other cognitive processes, and specifically that myopia for the future and the salience of jackpots control the propensity to take a gamble.

Results and Discussion

According to neoclassical economic theory, rational decision makers should be indifferent when offered a choice between a safe option and a risky option tendering the same average expected payoff [8–10]. Economists and psychologists have long known, however, that humans and animals deviate from this normative ideal. Typically, choosers will pay a premium for the safe option, although some individuals favor, and some contexts increase the likelihood of, choosing the risky alternative [7, 11–13]. Individuals may deviate from strict rationality because of nonlinear internal representation of risk [12, 14], nonlinear effects of payoffs on fitness

[15, 16], decreasing marginal utility of income [17], or Weber's law effects on discrimination of rewards [11, 18].

Just as individuals tend to systematically deviate from normative ideals when choosing between risky options, they also tend to undervalue, or discount, goods offered at future times and impulsively favor small but immediate gains [19–22]. Furthermore, time and risk both lead to hyperbolic discounting [3, 4]. The qualitative and quantitative similarities between risk sensitivity and temporal discounting suggest a common underlying mechanism [1–6]. Dysfunctional decision making in a number of neuropsychiatric disorders, such as addiction, problem gambling, and impulsive behavior, also supports the hypothesis that shared processes contribute to risk sensitivity and temporal discounting.

On a single trial of a standard gambling task, the risky option offers only a 50% chance of a large payoff, but if the subject adopts a strategy of choosing the risky option and sticks with it for a series of trials, the risky strategy offers a virtually guaranteed large payoff, albeit delayed. Thus, one way to understand the relationship between risk sensitivity and temporal discounting is to assume that future probabilistic rewards are construed as certain rewards and discounted over time in the same way as certain rewards [5, 6, 23]. This model predicts that propensity to gamble should decrease as delay between choices (typically determined by the intertrial interval, ITI, in most experimental contexts) increases. In a previous study, our lab demonstrated that rhesus macaques strongly prefer the risky option in a visual gambling task and, moreover, that risk seeking increases with increasing variance in payoff of the risky option [7]. In that study, monkeys made choices on average once every 3 s. Analogous studies in rats and birds, however, have presented choices with much longer delays between trials (normally 30 s) and found that animals were risk averse. Similarly, studies of human risk preferences often employ one-shot gambles—effectively an infinite ITI—and most people also avoid the risky option [11–13, 24]. These observations suggest that temporal discounting may in fact explain much of the variance in risk preference observed in different studies.

Based on these considerations, we hypothesized that risk preference in monkeys would vary systematically with the delay between choices in which risk remained constant. To test this hypothesis, we examined choices made by rhesus monkeys performing a visual gambling task (Figure 1) [7]. Monkeys sat in front of a computer monitor and fixated a central “start” stimulus. Two identical choice targets then appeared, one on the left and the other on the right. The “safe” target offered a guaranteed reward (150 ms solenoid open time of water) while the “risky” target offered an even gamble between a larger and a smaller reward (250 ms and 50 ms of water). To encourage sampling of both locations, the identity of the risky and safe targets was switched (with no overt signal to the subjects) every 25 trials. Critically,

*Correspondence: hayden@neuro.duke.edu

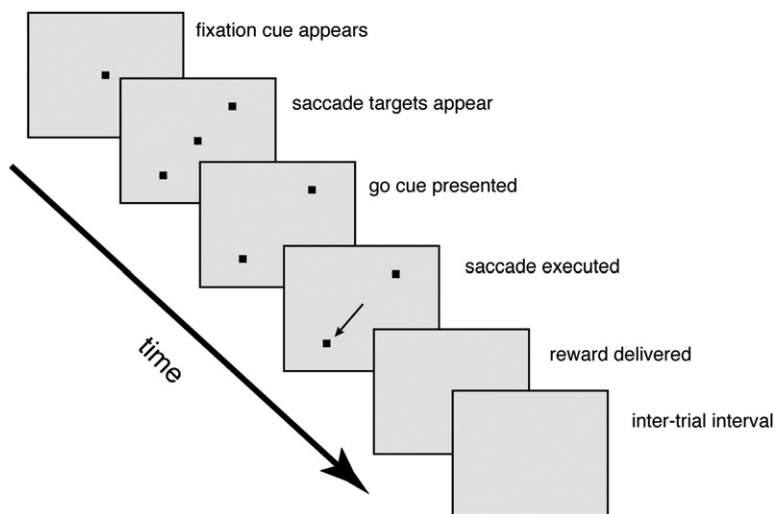


Figure 1. Structure of Visual Gambling Task
Each frame represents a stage of the trial. Trial begins when fixation cue (small central square) appears. Two response targets appear on opposite sides of the fixation spot. After a 1 s decision period, the central cue is extinguished, indicating that gaze must be shifted to either of the two targets within 100 ms. After choice, all stimuli are extinguished and reward is delivered. Reward was 150 ms of fluid for choice of the safe target and either 50 ms or 250 ms of fluid for the risky target. After reward, the screen remains blank for an inter-trial interval (ITI) of variable duration.

we varied the time between choices by changing the ITI randomly every 50–200 trials (ITIs: 1 s, 3 s, 6 s, 10 s, 12 s, 20 s, 30 s, 45 s, 60 s, 90 s). Monkeys typically performed 300–600 trials in a session.

We predicted that preference for the risky option would systematically diminish with increasing time between choices. Figure 2 shows the average preference for the risky option for both monkeys as a function of ITI. As the delay between trials increased, preference for the risky option systematically declined (logistic regression on raw choice data, $r = 0.0152$, $p < 0.0001$). The same inverse relationship between preference for the risky option and ITI was observed in both monkeys individually (monkey Niko, $r = 0.0105$, $p < 0.0001$; monkey Broome, $r = 0.0220$, $p < 0.0001$).




Figure 2. Increasing Time between Choices Reduces Preference for Risk in Monkeys

Probability of choosing the risky target plotted as a function of inter-trial interval (ITI) for two monkeys. As ITI becomes longer, risk preference declines. Open circles show mean values, solid dots show one standard error above and below the mean (calculated via bootstrap). Solid line indicates best-fit linear regression, and its slope is significantly negative ($p < 0.0001$). Dashed line indicates risk neutrality.

If monkeys construe the risky option as offering a virtually certain large reward at a randomly varying future time, then the subjective expected utility (SEU) of the risky option will depend on the time at which the higher payoff is expected. Because the expected large reward may occur on a future trial, the SEU of the risky option depends on the ITI. Thus, the hypothesis that monkeys construe their decision as offering a virtually certain large payoff at a variable future time leads to the hypothesis that ITI, a factor beyond the scope of the current trial, will influence the appeal of the risky option, and thus change the probability that it will be chosen [5, 6, 23].

To test this hypothesis, we estimated SEUs for risky options based on the expected delay until the payoff. We assumed a hyperbolic discount function [19, 25] with a discount factor (k) of 0.04, estimated from responses of rhesus macaques performing an intertemporal choice task for fluid rewards (K. Louie and P. Glimcher, 2006, Soc. Neurosci., abstract). We used this technique to estimate SEUs for the risky option at a range of delays (solid curve in Figure 3). Note that the SEU for the safe option was identical to its expected utility and did not depend on ITI (150 ms, dashed line). We next predicted behavioral choices from the calculated SEUs. Earlier research in our lab demonstrated that choices between two safe options depend on the ratio of their values; the relationship between reward ratio and behavioral choice was well approximated by a linear function within the limits of the model [26]. With this information, the predicted choice frequencies matched the observed risk preferences remarkably well (Figure 3). Moreover, the hyperbolic function gave a better fit to the data than a linear model (average error = 6.4 ms and $r^2 = 0.86$ for the hyperbolic model and 7.4 ms and $r^2 = 0.75$ for the linear model). We note that our monkeys remained risk neutral at the longest delay (90 s); we were unable to test values beyond 90 s because the monkeys tended to fall asleep during these longer intervals.

While students of animal behavior generally assume that the form of the temporal discounting function is hyperbolic, economists often describe it as exponential [1], and some behavioral studies have suggested that discounting may be exponential in monkeys as well (J. Hwang and D. Lee, 2005, Soc. Neurosci., abstract).




Figure 3. Temporal Discounting Predicts Risk Preferences in Monkeys

Subjective expected utility (SEU) of risky target (left ordinate) as a function of intertrial interval as estimated by the behavioral model assuming hyperbolic discounting (see [Experimental Procedures](#)). Probability of a risky choice is overlaid to facilitate comparison (dots, right ordinate). The model predicts that the SEU of the risky and safe targets will be equal with an ITI of 46 s. Behavioral data is well fit by the model ($r^2 = 0.86$, $p = 0.0001$). Note that the behavioral data show a small tendency to deviate from the model data at very short ITIs. Dotted line indicates predictions of best-fit exponential model (see text). Dashed line indicates risk neutrality, and also shows the predicted SEU of the safe target.

To test whether our results depend on the form of discounting we assumed, we predicted risk preferences with an exponential model of the form $v' = ve^{kD}$. Although the best-fit exponential model ($k = -0.019$, average error = 10.5 ms, $r^2 = 0.73$) was a poorer fit than the hyperbolic model, it did provide a qualitative match to the data. Our data therefore are consistent with both forms of temporal discounting.

To assess the validity of the assumed hyperbolic discount parameter, $k = 0.04$, we tested the goodness of fit of hyperbolic functions to the data assuming discount parameters ranging from 0 to 1. We found that the discount parameter that most closely fit the data was 0.033, very close to the assumed parameter of 0.04 (r^2 for predictions associated with this k -value is 0.87). The close agreement between this value and discount parameters estimated in other studies of intertemporal choice in primates (K. Louie and P. Glimcher, 2006, *Soc. Neurosci.*, abstract) [27] further supports this model. We also note that this method offers a new tool for calculating discount parameters without using an explicit intertemporal choice procedure.

One potential confound that could account for our results is that monkeys may forget the reward values associated with the two targets as ITI increases. To control for this possibility, we performed an additional experiment offering a choice between two close but different safe options (150 and 180 ms access to water). In all other details, the methods of this experiment were the same as those described in the experiment above. Monkeys strongly preferred the larger reward, and this preference did not depend on ITI (Figure 4). These results demonstrate that monkeys can remember the payoffs associated with each target for the longest durations




Figure 4. Effects of ITI on Risk Preference Does Not Reflect Forgetting

Mean probability 0.66 (standard deviation ± 0.05) of a choosing the larger reward for two monkeys plotted as a function of ITI in a control experiment with two safe options of different value. Solid dots show one standard error above and below the mean. There was no significant effect of ITI on choice.

we tested and that the observed decline in risk seeking with increasing ITI cannot be explained by forgetting.

It is worth making explicit that two specific deviations from rationality are inherent in the model we propose. Besides hyperbolic discounting, the model suggests that when evaluating the risky option, monkeys look forward only until the next large payoff [5]. In focusing attention on the jackpot, subjects overvalue the SEU of the risky option. This overvaluation of large payoffs results in risk-seeking behavior at shorter ITIs. There is no a priori reason why monkeys do not construe the risky option as a series of large payoffs followed by a small one; however, this would lead to increased risk seeking with increasing ITIs, a choice pattern not observed here.

More generally, it is worth making explicit that ours is not the only model that could explain the data presented here. Other models that could explain our results qualitatively include optimal foraging and state-dependant models of risk [16, 18, 28]. In contrast to the model proposed here, these models do not require subjective devaluation of the future and thus do not invoke the concept of subjective expected utility. Future research will be needed to distinguish these models from the proposed one.

Because human gamblers typically discount the future more steeply than normal adults, it might be predicted that they would find gambling more aversive than normals do. To resolve this well-known paradox, it has been proposed that an increased salience of jackpots drives gamblers to find gambling more rewarding than normals do, and that their steeper discount functions drive them to prefer the immediate smaller rewards of gambling to the longer rewards associated with saving [5].

Conclusions

We demonstrated that risk sensitivity in rhesus monkeys depends strongly on the time between subsequent

choices. These results indicate that contextual factors that appear to be unrelated to the task at hand can significantly influence reward-guided decision making. The few studies directly investigating the role of ITI on risk preference have obtained inconsistent results [6, 14, 23], possibly a consequence of the abstract nature of the monetary rewards used in those experiments. The present results illustrate that contextual influences on risk sensitivity are more general than previously thought and extend beyond the horizon of the current trial. Moreover, individual differences in risk sensitivity may simply reflect differences in the discount factor k , which could be specified genetically and manipulated pharmacologically, thus offering new potential therapies for problem gambling and addiction.

Experimental Procedures

Two adult male rhesus monkeys (*Macaca mulatta*) served as subjects. Subjects were on controlled access to fluid outside of experimental sessions; they earned roughly 80% of their total daily fluid ration during experimental sessions. All procedures were approved by the Duke University Medical Center Institutional Animal Care and Use Committee and complied with the Public Health Service's Guide for the Care and Use of Animals.

A standard desktop computer controlled experiments and recorded data with custom software (Gramalkyn, <http://ryklinsoftware.com>). Stimuli were presented on a 21" Sony Trinitron monitor (1024 × 768 resolution; 60 Hz refresh) placed directly in front of the monkey chair, 45 cm away. Stimuli were small yellow squares subtending approximately 1° of visual angle. All stimuli were presented on a dark background. Eye position was monitored at 500 Hz with a scleral search coil implanted with standard techniques described previously [29] or with an Eyelink 1000 Camera System (SR Research, Osgoode, ON, Canada).

A standard solenoid valve controlled the duration of water delivery. As in previous studies, we calibrated the water delivery system to ensure that reward volume was linearly proportional to valve open time [7]. Trials began when the monkey acquired fixation on a central square. When fixation had been achieved and maintained for 1 s, two response targets (yellow squares identical to the fixation cue) appeared at eccentric locations on opposite sides of the visual field. They were usually 15° to the left and right of the central square, although they were sometimes placed at 15° above and below the center, or at oblique orientations. The locations used did not affect choice.

After 1 s, the central fixation square was extinguished, and the monkey had 100 ms to shift gaze to either of the two targets. Monkeys were rewarded with water delivered from a solenoid valve. The volume of water was linearly proportional to solenoid open time, and 150 ms provided 0.023 ml of water [7]. One target offered a "safe" option (150 ms of water), while the other offered a "risky" option (50% chance of 50 ms of water and 50% chance of 250 ms of water). The smaller option was set at 50 ms so that monkeys would be rewarded in all conditions. Based on previous results, we hypothesized that had we used 300 ms and 0 ms as the two risky possibilities, behavior would have been qualitatively similar, but that risk seeking would have been stronger [7]. There was no overt cue indicating which target was risky and which was safe. To encourage foraging, the risky and safe sides were switched (with no cue given to the subjects) every 25 trials.

After delivery of the reward, all stimuli were removed, and the monitor was left blank for a specified duration (intertrial interval, ITI). ITIs were drawn randomly from a pool of values (1 s, 3 s, 6 s, 10 s, 12 s, 20 s, 30 s, 45 s, 60 s, 90 s) and varied in blocks of 50–200 trials. Monkeys were tested a total of 3860, 1943, 2291, 842, 607, 591, 1016, 112, 476, and 178 trials in each condition, respectively. To minimize the possibility that changes in observed risk preference on different days would cause spurious changes in risk preference, we presented blocks of each ITI in a random order, and we tried to present as many ITIs as possible each day. We were

able to present at least six blocks each day, and on half the days, we presented all ten. The length of each block was chosen arbitrarily, and the sequence of blocks was randomized. We also performed a control experiment in which both options offered safe payoffs. In this experiment, one target offered 150 ms of reward and the other offered 180. As in the other experiment, the identity of the two targets switched every 25 trials. In this experiment, we used only seven ITIs (1 s, 3 s, 6 s, 10 s, 30 s, 60 s, 90 s). In all experiments, monkeys typically earned around 300 ml of water in an experimental session.

The Model

We used a model assuming monkeys evaluate the discounted reward associated with the next large payoff to determine the subjective expected utility (SEU) of the risky option. That is, in the model, there is a 50% chance of an immediate large (250 ms) payoff and a 50% chance of a deferred large payoff. (The model explicitly assumes that the risky choice is evaluated as if the risky choice would be consistently chosen on all future trials.) On the deferred payoff trials, there is a 50% chance that a large payoff will occur after only one trial (plus ITI), and a 50% chance that the large payoff will be further delayed. This thought process may be extended out infinitely. To determine the SEU of the risky option, we calculated the first 100 terms of the infinite series. This process sums each possible outcome multiplied by its probability. Each outcome is defined as the average of the discounted payoffs for all trials up to and including the first large payoff. In practice, we assume the brain estimates the SEU of each option based on experienced rewards and the intervening intervals [30].

We calculated the discounted utilities with a hyperbolic discount function of the form $v' = v/(1 + kD)$, where v' refers to the discounted utility, v refers to the objective value (in ms of fluid), k refers to the discount parameter, and D refers to the delay (in seconds) until the payoff is offered [19]. We used a discount parameter k of 0.04 (K. Louie and P. Glimcher, 2006, Soc. Neurosci., abstract). The immediate payoff was assumed to have a delay of 0 s and the subsequent payoffs had a delay of $(3 + ITI)^n$, where n refers to the trial number in the future, and 3 refers to the average duration of the trial itself. We do not suppose that the monkeys explicitly follow our method to evaluate SEUs, but merely that they use cognitive strategies that produce the same results as this method.

To calculate the accuracy of the exponential discounting model, we used the standard exponential discounting equation, $v' = ve^{kD}$. We then obtained the best-fit parameters for this model, which provided a value of $k = -0.019$.

To predict choice probability as a function of SEU, we applied the results of earlier research performed in our lab showing a roughly linear relationship between the relative value of two options in a decision-making task and choice probability [26]. These data (collected on the same two monkeys used in the present experiment) show that in the range of values used in this experiment, choice probability is well approximated by a linear function of the ratio of the two options, specifically, $p = (\text{ratio} - 0.4) * 5$.

Supplemental Data

Supplemental Data include two figures and Supplemental Experimental Procedures and can be found with this article online at <http://www.current-biology.com/cgi/content/full/17/1/49/DC1/>.

Acknowledgments

We thank G. Madden (University of Kansas) for discussions that led to the development of this project. This work was supported by the National Institute of Health.

Received: September 26, 2006

Revised: October 26, 2006

Accepted: October 27, 2006

Published: January 8, 2007

References

1. Green, L., and Myerson, J. (1996). Exponential versus delayed outcomes: risk and waiting time. *Integr. Comp. Biol.* 36, 496–505.

2. Mischel, W., and Grusec, J. (1967). Waiting for rewards and punishments: effects of time and probability on choice. *J. Pers. Soc. Psychol.* 5, 24–31.
3. Green, L., and Myerson, J. (2004). A discounting framework for choice with delayed and probabilistic rewards. *Psychol. Bull.* 130, 769–792.
4. Prelec, D., and Loewenstein, G. (1991). Decision making over time and under uncertainty: a common approach. *Manage. Sci.* 37, 770–786.
5. Rachlin, H. (2000). *The Science of Self-Control* (Cambridge, MA: Harvard University Press).
6. Rachlin, H., Logue, A.W., Gibbon, J., and Frankel, M. (1986). Cognition and behavior in studies of choice. *Psychol. Rev.* 93, 33–45.
7. McCoy, A.N., and Platt, M.L. (2005). Risk-sensitive neurons in macaque posterior cingulate cortex. *Nat. Neurosci.* 8, 1220–1227.
8. Von Neumann, J.V., and Morgenstern, O. (1944). *Theory of Games and Economic Behavior* (Princeton, NJ: Princeton University Press).
9. Stephens, D.W., and Krebs, J.R. (1986). *Foraging Theory* (Princeton, NJ: Princeton University Press).
10. Charnov, E.L. (1976). Optimal foraging, the marginal value theorem. *Theor. Popul. Biol.* 9, 129–136.
11. Weber, E.U., Shafir, S., and Blais, A.R. (2004). Predicting risk sensitivity in humans and lower animals: risk as variance or coefficient of variation. *Psychol. Rev.* 111, 430–445.
12. Kahneman, D., and Tversky, A. (1979). Prospect theory: an analysis of decision under risk. *Econometrica* 47, 263–291.
13. Kuhnen, C.M., and Knutson, B. (2005). The neural basis of financial risk taking. *Neuron* 47, 763–770.
14. Silberberg, A., Murray, P., Christensen, J., and Asano, T. (1988). Choice in the repeated-gambles experiment. *J. Exp. Anal. Behav.* 50, 187–195.
15. Smallwood, P. (1996). An introduction to risk sensitivity: the use of Jensen's inequality to clarify evolutionary arguments of adaptation and constraint. *Am. Zool.* 36, 392–401.
16. Kacelnik, A., and Bateson, M. (1996). Risky theories—the effects of variance on foraging decisions. *Am. Zool.* 36, 402–434.
17. Glimcher, P.W., and Rustichini, A. (2004). Neuroeconomics: the consilience of brain and decision. *Science* 306, 447–452.
18. Kacelnik, A., and Brito e Abreu, F. (1998). Risky choice and Weber's Law. *J. Theor. Biol.* 194, 289–298.
19. Mazur, J.E. (1987). An adjusting procedure for studying delayed reinforcement. In *Quantitative Analyses of Behavior, Volume 5. The Effect of Delay and Intervening Events on Reinforcement Value*, M.L. Commons, J.E. Mazur, J.A. Nevin, and H. Rachlin, eds. (Mahway, NJ: Erlbaum), pp. 55–73.
20. Bickel, W.K., and Marsch, L.A. (2001). Toward a behavioral economic understanding of drug dependence: delay discounting processes. *Addiction* 96, 73–86.
21. McClure, S.M., Laibson, D.I., Loewenstein, G., and Cohen, J.D. (2004). Separate neural systems value immediate and delayed monetary rewards. *Science* 306, 503–507.
22. Frederick, S., Loewenstein, G., and O'Donoghue, T. (2002). Time discounting and time preference: a critical review. *J. Econ. Lit.* 40, 351–401.
23. Rachlin, H., Raineri, A., and Cross, D. (1991). Subjective probability and delay. *J. Exp. Anal. Behav.* 55, 233–244.
24. Holt, C.A., and Laury, S.K. (2002). Risk aversion and incentive effects. *Am. Econ. Rev.* 92, 1644–1655.
25. Kirby, K.N., and Marakovic, N.N. (1996). Delay-discounting probabilistic rewards: rates decrease as amounts increase. *Psychon. Bull. Rev.* 3, 100–104.
26. McCoy, A.N., Crowley, J.C., Haghghian, G., Dean, H.L., and Platt, M.L. (2003). Saccade reward signals in posterior cingulate cortex. *Neuron* 40, 1031–1040.
27. Stevens, J.R., Hallinan, E.V., and Hauser, M.D. (2005). The ecology and evolution of patience in two New World primates. *Biol. Lett.* 1, 223–226.
28. Marsh, B., Schuck-Paim, C., and Kacelnik, A. (2004). Energetic state during learning affects foraging choices in starlings. *Behav. Ecol.* 15, 396–399.
29. Platt, M.L., and Glimcher, P.W. (1997). Responses of intraparietal neurons to saccadic targets and visual distractors. *J. Neurophysiol.* 78, 1574–1589.
30. Sugrue, L.P., Corrado, G.S., and Newsome, W.T. (2004). Matching behavior and the representation of value in the parietal cortex. *Science* 304, 1782–1786.