Journal of Experimental Psychology: Applied
2006, Vol. 12, No. 3, 129–141

1076-898X/06/$12.00 DOI: 10.1037/1076-898X.12.3.129

Risk Attitude in Small Timesaving Decisions

Nira Munichor and Ido Erev
Technion—Israel Institute of Technology

Arnon Lotem
Tel-Aviv University

Four experiments are presented that explore situations in which a decision maker has to rely on personal experience in an attempt to minimize delays. Experiment 1 shows that risk-attitude in these timesaving decisions is similar to risk-attitude in money-related decisions from experience: A risky prospect is more attractive than a safer prospect with the same expected value only when it leads to a better outcome most of the time. Experiment 2 highlights a boundary condition: It suggests that a difficulty in ranking the relevant delays moves behavior toward random choice. Experiments 3 and 4 show that when actions must be taken during the delay (thereby helping compare delays), this increases the similarity of timesaving decisions to money-related decisions. In these settings the results reflect an increase in risk aversion with experience. The relationship of the results to the study of non-human time-related decisions, human money-related decisions and human time perception is discussed.

Keywords: prospect theory, small decisions, payoff variability effect, hot-stove effect, RELACS

Many natural activities involve small timesaving decisions. During these activities the actor (a decision maker) selects among alternative courses of action in an attempt to save time. For example, consider the task of taking off one's shoe. During this activity one can choose between untying the laces first or trying to take off the shoe while the laces are still tied. The latter option may eliminate two or three steps (bending, untying, etc.), thereby saving a few moments. But if the shoe gets stuck, this alternative can increase the number of necessary steps and thus end up taking longer. In another example, consider the task of opening a document using Microsoft Windows. The actor can try to save time by opening the list of recent documents first. However, this method is likely to be counterproductive if the document is not in that list.

Small timesaving decisions of the type exemplified above do not seem very important. This is especially true for cases in which the difference between the outcomes involves only a few seconds. Certainly, there is no reason to put serious thought into each such small choice. Nevertheless, on the aggregate, small timesaving decisions can be highly consequential. For example, timesaving decisions can determine the effectiveness of practicing basic skills. As shown by Siegler and Lemaire (1997), the choice among alternative strategies for solving math problems is guided by the time (on the order of seconds) that each strategy takes. Computers offer another example: Gray and Boehm-Davis (2000) demonstrated that people can be sensitive to a 150-ms difference between strategies for moving the cursor and clicking on a button that appears on a computer screen. This and similar sensitivities bear important implications for the designers of interactive technologies: Differences on the order of seconds and even milliseconds can affect choices among software packages and search engines.

The current article focuses on pure timesaving decisions that are made repeatedly on the basis of personal experience. The focus on pure timesaving decisions implies that in the situations considered, minimizing delay is the sole goal of the decision maker; there is no tradeoff between this and other goals. The decision of how to take off a shoe, described above, is an example of this type of decision. First, it is a personal decision made daily, on a repeated basis. Second, the outcome of each alternative action involves time currency per se (money, for example, does not play a role here). Previous studies on the effect of time on decision making have usually focused on the tradeoff between time and other outcomes (see, e.g., Chapman et al., 2001; Chapman, Nelson, & Hier, 1999; Kirby & Herrnstein, 1995; Klatzky, 2000; Lowenstein & Elster, 1992; Lowenstein & Thaler, 1989). Hence, the present focus on pure timesaving decisions distinguishes the current study from most previous studies on the effect of time on decision making. In addition, the present focus on decisions that are made on the basis of personal experience implies that in the situations considered here, the decision maker cannot take advantage of precise descriptions of the possible distributions of delays but rather accumulates knowledge about possible outcomes through experience. Thus, the noisy nature of time perception (see Grondin, 2001; Hornik & Zakay, 1994; James, 1950) is likely to affect behavior.

The main goal of the current research is to improve understanding of the basic properties of small timesaving decisions and their relationship to the known properties of money-related decisions. Specifically, we focus on risk attitude in timesaving decisions that are made based on personal experience. We begin with a review of previous research, highlighting three reasonable but contradictory predictions regarding risk attitude in small timesaving decisions. A generalization from scenario-based studies predicts risk aversion,

Nira Munichor and Ido Erev, Max Wertheimer Minerva Center for Cognitive Studies, Behavioral Sciences, Faculty of Industrial Engineering and Management, Technion—Israel Institute of Technology, Haifa, Israel; Arnon Lotem, Department of Zoology, Faculty of Life Sciences, Tel-Aviv University, Tel-Aviv, Israel.

This research was supported by a grant from the Israel Science Foundation.

We thank Ralph Hertwig and Meira Ben-Gad for their helpful comments on earlier versions of this article.

Correspondence concerning this article should be addressed to Ido Erev, Max Wertheimer Minerva Center for Cognitive Studies, Behavioral Sciences, Faculty of Industrial Engineering and Management, Technion—Israel Institute of Technology, Haifa 32000, Israel. E-mail: erev@tx.technion.ac.il

whereas a generalization from studies of animal foraging behavior predicts risk seeking. A generalization from studies of small money-related decisions, however, predicts a tendency to prefer the alternative that minimizes delays in most cases. Four experiments that explore these hypotheses are presented.

Previous Research and Three Contradictory Predictions

Our search for previous studies that could be used to derive an unambiguous prediction for timesaving decisions led us to consider three lines of research. First, we looked at scenario-based studies of choice among distributions of time delays (see Krishnamurthy & Kumar, 2002; Leclerc, Schmitt, & Dubé, 1995). In most cases, this line of research demonstrates a bias toward risk aversion in time-related decisions under uncertainty. For example, in a series of studies, Leclerc et al. (1995) presented their participants with various scenarios that involved a selection between two alternatives carrying an identical expected time loss. In each scenario (waiting for a bus, waiting for a flight, or taking a ride) participants were required to choose between a certain time loss and a variable loss. The certain loss was selected by 70% to 86% of the participants. In other words, participants demonstrated a tendency to avoid risk for timesaving decisions, making risk-averse choices. A generalization of these results to the current setting, which also involves time-related decisions, predicts risk aversion in experience-based timesaving decisions.

A second line of relevant research focuses on nonhuman foraging behavior. The findings indicate that animals tend to be risk prone in decisions among delays in obtaining rewards and prefer a variable delay to a fixed delay (for a review see Kacelnik & Bateson, 1996). For example, Bateson and Kacelnik (1995) gave their starling subjects a choice between two keys, where one delivered food after a fixed delay of 20 s and the other, after a delay of 2.5 s or 60.5 s with equal probability. All 6 subjects demonstrated almost exclusive preference for the variable delay, choosing the variable option in 97.7% of trials. Kacelnik and Brito e Abreu (1998) noted that results of this kind are consistent with the predictions of prospect theory (Kahneman & Tversky, 1979), because in both cases subjects prefer the variable option when the outcomes are undesirable (involve a loss of money or time). To explain this attitude, they use a process-based model, known as scalar utility theory, which is based on the combined effect of perceptual error and Weber–Fechner's law (for a review of Weber's law for time intervals, see Gibbon, 1991). A generalization of these findings to the current context seems reasonable in light of the recently discovered similarity between human and nonhuman risk-taking behavior in experience-based decisions (Weber, Shafir, & Blais, 2004). Such a generalization would predict risk seeking in small timesaving decisions.

A third line of research focuses on small money-related decisions made by people. A generalization of this research to the current context leads to two related predictions. The first involves the effect of rare (low-probability) events. Study of money-related decisions demonstrates that when people rely on previous personal experience they behave as if they underweight rare events, tending to prefer the alternative that leads to a better outcome most of the time even when this alternative is not associated with a higher expected outcome. For example, in each of the 200 trials in the study of Barron and Erev (2003), participants were asked to choose between two keys, one leading to a loss of 9 points with certainty, and the other to a 10-point loss in 90% of the trials and 0 points in the rest. The results of this study reveal a tendency toward the first, "safer" key (a loss of 9 points with certainty), which most of the time (90% of the time) would result in a better outcome (a loss of 9) than the other key (a loss of 10). The "safer" key was selected in 63% of the trials. Barron and Erev suggested that this finding could be the result of a bias toward probability matching (see Estes, 1950) or of a tendency to rely on recent outcomes. As rare events (the 10% outcome) are less likely to have occurred recently, they are underweighted. Under this interpretation, it is natural to expect a similar pattern in time-related decisions, that is, a contingent risk attitude is predicted: Risk aversion is predicted when the safer prospect leads to a better perceived outcome most of the time, and risk seeking is predicted when the "risky" prospect leads to a better perceived outcome most of the time.

In addition, a generalization from the study of money-related decisions implies sensitivity to available feedback. In particular, this research demonstrates that when the feedback is limited to the *obtained outcome only* (the decision maker does not receive information concerning the forgone payoff), risk aversion increases with experience. This pattern is explained by the "hot-stove" (or "stickiness") effect (see Denrell & March, 2001; Erev & Barron, 2005). The explanation is based on the assumption that bad experiences decrease the tendency to select the same alternative again. Thus, when learning is based on limited feedback, bad experiences have a larger effect than good ones: They remain the most recent information about the relevant alternative for a longer time. As a result, the attractiveness of high-variability (risky) alternatives decreases with experience.

Experiment 1: Time-Related Decisions With Complete Feedback

To clarify the differences between natural generalizations of the three lines of research considered above, Experiment 1 focused on the following choice problems:

<u>Problem 1</u>

Choose between:

S: Loss of 2.8 with certainty

R: Loss of 3 with probability 0.90; loss of 1 otherwise.

<u>Problem 2</u>

Choose between:

S: Loss of 3 with certainty

R: Loss of 2 with probability 0.875; loss of 10 otherwise.

In both problems the alternatives entailed time delays (in seconds), which were used as the immediate feedback. The participants received complete feedback (including the obtained and forgone payoff after each trial). At the conclusion of the experiment, delays were converted to monetary payoffs according to the following exchange rate: 1 s in Problem 1 or 3 s in Problem 2 equaled 1 Sheqel (approximately 22 U.S. cents). Different show-up fees were used to ensure similar expected payoffs.

Note that the two alternatives for each choice problem have an identical expected loss (of 2.8 and 3, respectively). The problems differ in their payoff matrices. In Problem 1 the alternative that leads to a better outcome most of the time is the "safer" alternative.

Most often, a participant who chose not to go with the certain outcome—a loss of 2.8 seconds—would end up losing 3 s, and thus perform worse than if she had stuck with the certain loss. In contrast, in Problem 2 the alternative that leads to a better outcome most of the time is the "risky" alternative. Here, going with the variable outcome would most often lead to better performance than sticking with the certain loss.

A generalization of the scenario-based research suggests a tendency to favor S over R in both problems. A generalization of the research on nonhuman foraging behavior, which suggests risk aversion in the loss domain, seems to imply the opposite pattern (although, as we discuss later, a careful examination of scalar utility theory suggests that it may not predict risk seeking for all cases of variable versus fixed delays). Finally, a direct generalization of the results observed in money-related decisions from experience suggests a tendency to prefer S in Problem 1 but R in Problem 2, because these alternatives have a better payoff most of the time. (The hot-stove effect is not relevant here as the feedback is complete. This effect is examined in Experiments 2, 3, and 4.)

Method

Participants. For all four experiments, participants were recruited through campus advertisements from a technological institute of learning in Israel. In Experiment 1, 54 Caucasian students (26 male and 28 female) served as paid participants in the experiment. Most of the participants were 2nd- and 3rd-year industrial engineering and economics majors who had taken at least one probability and economics course. They were randomly assigned to either "Problem 1 group" ($n = 24$) or "Problem 2 group" ($n = 30$). Participants received an initial fee of 300 Sheqels (approximately $67) for the Problem 1 group and 135 Sheqels (approximately $30) for the Problem 2 group, from which a performance-contingent payoff was subtracted. Final payoffs ranged between 14 and 53 Sheqels (approximately $3 to $12).

Apparatus and procedure. Participants were informed that they were playing a "two-button machine" (see instructions in Appendix A) but received no prior information as to the game's payoff structure. Their basic task was to select one of two unmarked buttons presented on a computer screen (see Figure B1, in Appendix B), which represented the alternatives for the relevant problem. This basic task was performed 100 times. To avoid an "end-of-task" effect (i.e., change in performance as the end of the task is approached, a time when participants are expected to maximize expected utility; Catalano, 1973), the experiment was designed to avoid informing participants that the experiment consisted of exactly 100 trials.

Payoffs were contingent on the button chosen by the participants and were drawn from the distribution associated with the selected button, described above (distributions were randomly assigned to buttons). The payoff for the choice and forgone payoff information appeared following each choice and served as immediate feedback. Specifically, the feedback mode involved:

1. An appearance of a red light on the selected key (surrounded by a black frame) for a duration equaling the specific loss of time; and
2. A simultaneous appearance of a red light on the nonselected key for the amount of time that the participant would have lost if he or she had selected this key.

Results and Discussion

The left-hand column of Figure 1 presents a graph of the proportions of risky choices (R) in 10 blocks of 10 trials over participants (the right-hand column shows the predictions of a model discussed below). Over trials the proportion of risky choices (preferring the gamble to the certain outcome) in Problem 1 was 0.32 ($SD = 0.24$). The distance from random choice (the effect) was moderate ($d = 0.75$) and significant, $t(23) = 3.74$, $p < .01$. The proportion of risky choices in Problem 2 was 0.63 ($SD = 0.19$). The distance from random choice (the effect) was moderate ($d = 0.68$) and significant, $t(29) = 3.62$, $p < .01$. The difference between the two proportions was large ($d = 1.43$) and significant, $t(52) = 5.27$, $p < .01$.

In addition, we applied a model with repeated measures for data analysis to examined whether the proportion of risky choices differed between blocks in each problem. The effect of the interaction between problems and blocks was significant, $F(9, 468) = 5.36$, $MSE = 0.10$, $p < .01$, suggesting that some learning continued to occur across trials.

The results depicted above are consistent with the pattern observed by Barron and Erev (2003) in their study of money-related decisions from experience. As predicted by a generalization of Barron and Erev's assertions to the current setting, the observed risk attitude can be described as a tendency to select the option that

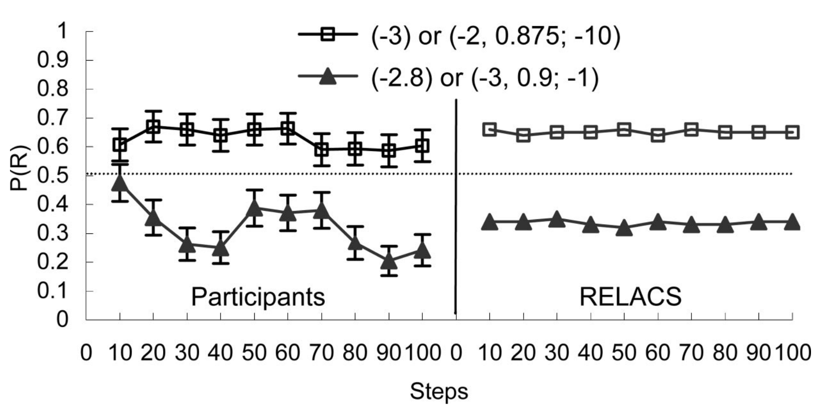

Figure 1. Experiment 1: The left-hand column presents the proportion of risky choices in Problem 1 (−2.8 vs. –3, 0.9; −1) and Problem 2 (−3 vs. –2, 0.875; −10) with complete feedback. The right-hand column presents the predicted proportion of risky choices according to the reinforcement learning among cognitive strategies (RELACS) model. Points represent the mean proportion over 10 trials; vertical lines depict 95% confidence intervals.

leads to a better outcome most of the time. Risk aversion was observed in Problem 1, where the "safer" option led to shorter delays most of the time (90% of the trials). However, in Problem 2, where the "safer" option led to longer delays most of the time (87.5% of the trials), participants exhibited risk seeking. Barron and Erev noted that their results imply underweighting of low-probability events. The current results suggest that the tendency to underweight rare events in repeated decisions from experience is not limited to decisions with numerical feedback. Here, a choice pattern that indicates a tendency to underweight rare events was also observed when the feedback consisted of time delays.

Experiment 2: Time-Related Decisions With Incomplete Feedback

Experiment 1 examined situations in which the decision maker received complete feedback that included information about both the obtained and forgone outcomes. Experiment 2, which focused on Problem 1, was designed to examine situations in which feedback was limited to the obtained outcomes; the actors did not receive feedback concerning the outcome of the option that they did not select. As noted above, study of money-related decisions predicts a hot-stove (stickiness) effect in such situations: When information is limited, experience is expected to increase risk aversion (Denrell & March, 2001; Erev & Barron, 2005).

There are, however, reasons to question the generalization of the hot-stove effect to timesaving decisions. Since absolute time perception is known to be noisy (see Hornik & Zakay, 1994; James, 1950), eliminating feedback concerning forgone payoffs is expected to impair the accuracy of how observed outcomes are ranked. When it is difficult to compare the delays associated with the different alternatives, different patterns of risk-attitude may emerge.

An attempt to derive the precise effect of perceptual noise in this setting reveals two different reasonable predictions. One prediction can be derived from a generalization of research showing the descriptive value of the lexicographic semi-order decision rule (Tversky, 1969; see related observations in Reyna & Brainerd, 1995). The current generalization is based on two working assumptions. The first of these implies that the choice alternatives in Problem 1 have two important dimensions: the *frequent outcome* (a loss of 3 s in the case of R and 2.8 s in the case of S) and the *best outcome* (a loss of 1 s for R, and 2.8 s for S). The second implies that elimination of the forgone payoff information reduces the probability that participants can distinguish between the two alternatives along the first dimension (the frequent outcomes of 2.8 s or 3 s). The lexicographic semi-order decision rule assumes that when alternatives cannot be reliably ranked along a particular dimension they are treated as being equal on that dimension (and the decision is made on the basis of different dimensions). Thus, a difficulty in distinguishing among frequent outcomes will increase the tendency to select the alternative with the higher best outcome (Alternative R).

A second prediction can be derived under the assumption that eliminating the forgone payoff information increases the perceived payoff variability, which in turn is assumed to decrease sensitivity to all payoff differences. As a result, the decision maker is likely to find it difficult to distinguish between the different options, and behavior is expected to move toward random choice (Myers & Sadler, 1960).

Method

Participants. The participants, none of whom had participated in Experiment 1, were recruited in the same manner as in Experiment 1. A total of 24 Caucasian (13 male, 11 female) university students participated in Experiment 2.

Apparatus and procedure. The materials and procedures were the same as in Experiment 1, except that feedback was limited to the obtained outcome.

Results and Discussion

The left-hand column of Figure 2 presents a graph of the proportions of risky choices (R) in 10 blocks of 10 trials over the participants. Across trials and participants, the proportion of risky choices was 0.55 (SD = 0.23). The distance from random choice (the effect) was small (d = 0.22) and not significant, $t(23) = 1.02$, *ns*. In addition, the results of a model with repeated measures do not reveal a consistent block effect, $F(9, 207) = 1.66$, MSE = 0.331, *ns*.

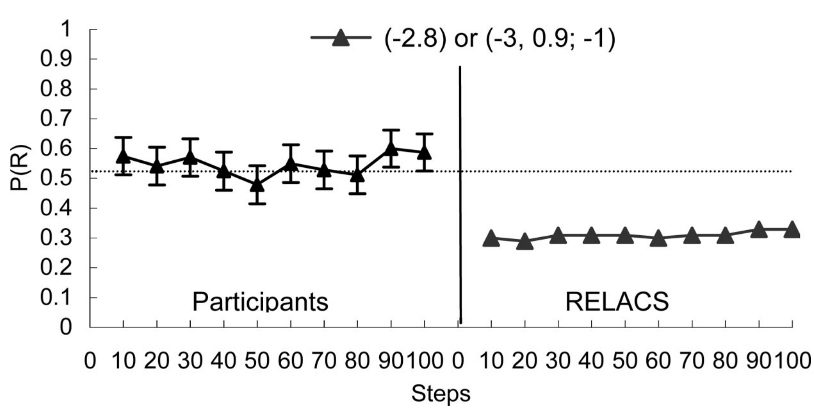

Figure 2. Experiment 2: The left-hand column presents the proportion of risky choices in Problem 1 (−2.8 vs. –3, 0.9; −1) with incomplete feedback. The right-hand column presents the predicted proportion of risky choices according to the reinforcement learning among cognitive strategies (RELACS) model. Points represent the mean proportion over 10 trials; vertical lines depict 95% confidence intervals.

The difference between the proportion of risky choices in Problem 1 without forgone payoff information (current condition) and the proportion with forgone payoff information (Problem 1 group in Experiment 1) was large ($d = 0.98$) and significant, $t(46) = 3.41$, $p < .01$.

These results are consistent with the assertion that elimination of the forgone payoff information increases the perceived payoff variability and moves choice behavior toward random choice. Note that in the current setting there is no significant tendency toward increased risk aversion with experience, as predicted by the hot-stove effect.

Experiment 3: Time- and Action-Saving Decisions

Experiment 3 was designed to explore the robustness of the results presented above in choice problems characterized by another two common properties of natural timesaving decisions. First, in natural timesaving decisions, payoffs are received in the currency of time. In Experiments 1 and 2, each second of delay was translated by the experimenter to a monetary loss. In Experiment 3, in contrast, participants were paid only for completing the experimental tasks, and were told that their goal was to complete them as swiftly as possible. Thus, each second they spent in the lab constituted a loss of one second of free time.

The second property involves the possibility that delays may not represent idle time but rather may entail performing additional actions to complete a task. In many natural examples of timesaving decisions, a decision maker must carry out a sequence of actions to complete a time-consuming task. For example, there are several ways to open a document using Microsoft Windows, in which an actor has to press different keys. In this and similar examples, timesaving implies minimizing the number of necessary actions or steps. Thus, the decision task in these settings may be perceived by the actor as an attempt to minimize the number of necessary steps.

Analyzing the expected effect of these properties leads to two contradictory hypotheses. The first is based on the working assumption that when discrete actions are required to complete a task, a decrease in the effect of perceptual noise can be expected. In simple terms, when the number of actions can be counted, it is easy to rank the different alternatives. As a result, "time- and action-saving decisions" are expected to be similar to money-related decisions even when the actors do not receive information concerning forgone payoffs: They are expected to reflect underweighting of rare events and the hot-stove effect.

The second hypothesis is more direct. The fact that the payoff involves only brief time delays makes the stakes extremely small (in Experiment 3 participants could lose time on the order of seconds only and risked no monetary loss). Thus, people are likely to be fairly indifferent between the different options. As a result, behavior in such settings is predicted to be close to random choice.

Experiment 3 was designed to explore these predictions. Two problems were compared:

Problem 2′ (a variant of Problem 2 from Experiment 1)

Choose between:

S: Loss of 3 s with certainty

R: Loss of 2 s with probability 0.875; loss of 10 s otherwise.

Problem 3

Choose between:

S: Loss of 3 s with certainty

R: Loss of 4 s with probability 2/3; loss of 1 s otherwise.

Note that although the two problems have identical expected losses (of 3 s), the implied distributions are quite distinct. In Problem 3 the safe alternative leads to a better outcome most of the time. In Problem 2′, however, the risky alternative leads to a better outcome most of the time. Thus, similarity to money-related decisions implies a higher level of risk-seeking choices in Problem 2′.

Method

Participants. Thirty-four Caucasian university students (19 male, 15 female), who did not participate in other experiments in this study, served as paid participants in the experiment. A within-subject design was used; each participant was faced with both problems in random order. The participants received only a show-up fee of 25 Sheqels (approximately $6.25) and could leave the laboratory immediately upon completing the experiment.

Apparatus and procedure. Participants received instructions informing them of the "game rules" (see instructions in Appendix C) but were given no prior information as to the game's payoff structure. Their basic task was to select one of two different-colored paths (representing the alternatives of the problem), whose starting points were presented on a computer screen (see Figure D1 in Appendix D). The selection initiated exposure of the chosen path, composed of squares. In each trial, the number of squares was determined by the outcome of the chosen alternative (e.g., if the outcome equaled 4, the path consisted of 4 squares). Participants were asked to press the "Enter" key once for each square that appeared. Each new square appeared 1 s after the "Enter" key was pressed. When the participant reached the end of the path, the next trial began. This basic task was performed 100 times for each problem. To avoid an end-of-task effect, participants were not informed that each part of the experiment consisted of exactly 100 trials.

Results and Discussion

The left-hand column of Figure 3 presents a graph of the proportions of risky choices (R) in 10 blocks of 10 trials over participants in each problem. Across trials the proportion of risky choices in Problem 2′ was 0.52 ($SD = 0.27$). The distance from random choice (the effect) was small ($d = 0.07$) and not significant, $t(33) = 0.39$, *ns*.The proportion of risky choices in Problem 3 was 0.36 ($SD = 0.21$). The distance from random choice (the effect) was moderate ($d = 0.67$) and significant, $t(33) = 3.99$, $p < .01$. Similarly, the difference between the two proportions was moderate ($d = 0.66$) and significant, $t(33) = 3.04$, $p < .01$. According to the results of a model with repeated measures, there was no consistent block effect, $F(9, 297) = 1.19$, $MSE = 0.101$, *ns*.

The results reflect the two properties of money-related decisions considered above. First, over the two conditions, participants appeared to prefer the alternative that leads to a better outcome most of the time. Second, the learning curves (cf. the left-hand column in Figure 3) exhibit the hot-stove effect: a decrease in risk seeking with time in the two conditions.

In summary, the results of Experiments 1, 2, and 3 can be summarized with the assertion that timesaving decisions are similar to money-related decisions when actors can accurately rank the attractiveness of the possible outcomes (Experiments 1 and 3) and are closer to random choice when reliable ranking is difficult (Experiment 2). To evaluate this verbal summary of the results, we compared the current findings with the predictions of the model proposed by Erev and Barron (2005) to capture behavior in

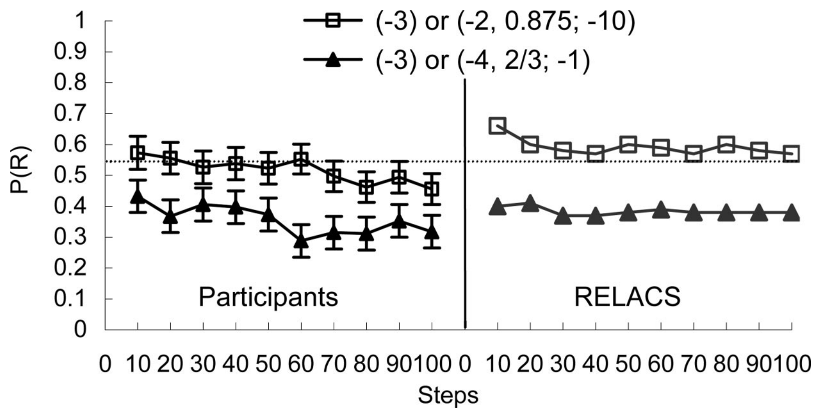

Figure 3. Experiment 3: The left-hand column presents the proportion of risky choices in Problem 2 (−3 vs. –2, 0.875; −10) and Problem 3 (−3 vs. –4, 2/3; −1) with incomplete feedback when decisions concern discrete time. The right-hand column presents the predicted proportion of risky choices according to the reinforcement learning among cognitive strategies (RELACS) model. Points represent the mean proportion over 10 trials; vertical lines depict 95% confidence intervals.

money-related decisions. According to this model, referred to as "reinforcement learning among cognitive strategies" (RELACS; see basic assumptions in Appendix E), decision makers learn through three basic cognitive strategies (choice rules). The first rule, "slow best reply," implies slow learning toward the preference of the alternative with the highest average payoff. The second strategy, "fast best reply," implies reliance on recent outcomes, and the third strategy, "case-based reasoning," implies (in this case) probability matching (see Estes, 1950).

The right-hand columns of Figures 1 through 3 show the predictions of the RELACS model for the proportion of risky choices in Experiments 1 through 3. The predictions were derived with the use of a computer simulation in which virtual agents that behaved according to the model's assumptions (with the parameters estimated by Erev & Barron, 2005) participated in a virtual replication of each of the five conditions. The results show a large difference between Experiment 2 and the other experiments. The mean squared distance (MSD score) between the experimental curve and the RELACS predictions was .00035 and .0054 in Experiment 1 and 3, respectively. These values are similar to the MSD scores found by Erev and Barron (2005) in validation of the model in money-related decisions. The MSD score in Experiment 2 was .0560, suggesting a poorer fit with the model's predictions under this condition. This result might have been expected, considering that ranking the relevant delays was much more difficult in Experiment 2 (as described above).

Experiment 4: Cross-Validation

Experiment 4 was designed to test the validity of the summary presented above. It examined 10 different experimental problems under two conditions: "red light" and "actions." The red light condition used the paradigm of Experiment 2: participants were asked to minimize the time consumed by presentation of a red light without receiving information about forgone payoffs (alternative delays). The actions condition followed the paradigm of Experiment 3: participants were asked to minimize the number of time-consuming actions required to complete a task. The problems were randomly selected from a space of possible conditions. The left-hand column in Table 1 presents the 10 problems and the algorithm used to generate them (sample them from the relevant space).

Method

Participants. Forty Caucasian students (23 male, 17 female) served as paid participants in the experiment. None of the participants took part in any of the previous experiments in this study. Among the sample, 20 students were assigned to each of the two conditions. Each participant received a show-up fee of 50 Sheqels (approximately $12.50) and could leave the laboratory immediately upon completing the experiment.

Apparatus and procedure. The apparatus and procedures were the same as in Experiment 2 (in the red light condition) and Experiment 3 (in the actions condition), with the exception that each participant faced 10 problems and each problem included 50 trials.

Results and Discussion

The right-hand column in Table 1 shows the observed and predicted proportions of risky choices over the 50 trials in each problem. Figure 4 presents the observed and predicted learning curves. The results support the summary of Experiments 1, 2, and 3 presented above: In the actions condition, the results reflected the basic properties of money-related decisions as captured by RELACS. Indeed, in all eight problems in which RELACS predicts deviation from risk neutrality (and from random choice), the mean results exhibited deviation in the predicted direction. The correlation between the observed and predicted proportion of risky choices was $r(18) = .71$, $p < .01$. The MSD between the observed and predicted proportion of risky choices was 0.0053. This value was similar to the values observed above and lower than the MSD between the observed proportion of risky choices and the prediction of a model that assumes risk neutrality (random choice; 0.012). In the red light condition, the observed behavior was closer to random choice (MSD of 0.0042) than to the predictions of RELACS (MSD of 0.0046). The correlation between the observed and predicted proportion of risky choices by RELACS was $r(18) = 0.42$, *ns.*

Table 1
The 10 Problems Produced by the Algorithm and the Predicted and Observed Average Proportion of Risky Choices in Experiment 4

	The 10 problems[a]						Observed and predicted proportion of risky choices		
Problem no.	$X(2)$	p	$X(3)$	$X(1)$	q	$X(4)$	The actions condition	The red light condition	RELACS model
1	−2	0.5	−3	−2	0.83333	−5	0.52	0.49	0.53
2	−2	0.5	−3	−1	0.50000	−4	0.31	0.43	0.46
3	−3	0.5	−4	−1	0.37500	−5	0.27	0.45	0.41
4	−3	0.5	−6	−2	0.37500	−6	0.44	0.63	0.50
5	−5	0.5	−6	−3	0.16667	−6	0.43	0.48	0.41
6	−1	0.5	−2	−1	0.87500	−5	0.51	0.54	0.54
7	−1	0.5	−4	−1	0.70000	−6	0.46	0.55	0.50
8	−3	0.5	−5	−2	0.50000	−6	0.48	0.41	0.47
9	−2	0.5	−3	−1	0.25000	−3	0.45	0.57	0.46
10	−3	0.5	−4	−2	0.50000	−5	0.40	0.51	0.46

Note. RELACS = reinforcement learning among cognitive strategies.
[a] The 10 problems were selected using the following four-step algorithm:
1. Four values, $X(i)$; $i = 1,2,3,4$, were randomly selected (without replacement) from the set: 1,2,3,4,5,6, and were ranked in descending order ($X(1)$ was the highest value).
2. The two potential prospects were set to equal:
   S: Loss of $X(2)$ with probability 1/2, loss of $X(3)$ otherwise.
   R: Loss of $X(1)$ with probability q, loss of $X(4)$ otherwise.
3. The value of q was computed to ensure that the two prospects had the same expected value.
4. The new problem was added to the experiment only if $X(2) \neq X(3)$, and $X(1) \neq X(2)$ or $X(3) \neq X(4)$.

To evaluate the importance of the difference between the two conditions, we focused on the difference between the MSD score of a model that assumes random choice and that of RELACS. The mean magnitude of this difference was 0.006 ($SD = 0.01$) in the actions condition and −0.0005 ($SD = 0.004$) in the red light condition. The implied effect was large ($d = 0.85$), $t(9) = 1.94$, $p < .05$, in a one-tailed test.

Additional analysis of the learning curves shows that the data in the actions condition reflect the hot-stove effect predicted by RELACS. The proportion of risky choices decreased with experience in all 10 problems: $P(\mathrm{R})$ in the last 20 trials is lower than $P(\mathrm{R})$ in the first 20 trials. Evaluation of the learning curves in the red light condition does not reveal a clear pattern. The proportion of risky choices decreased with experience in 6 problems and increased with experience in 4 problems.

It is important to note that whereas RELACS provides useful predictions of behavior in the actions condition, this does not mean that these predictions are accurate. Indeed, it seems that the predictions are biased: In 9 of the 10 problems in the actions condition the observed results reflected a lower proportion of risky choices than the RELACS predictions. Under one interpretation of this observation, the results reflect a bias toward consistency seeking. In repeated action-saving and timesaving decisions, the safer alternative maximizes the consistency between the different actions that must be taken in the different trials. Given that consistent mapping is known to reduce reaction time (Schneider & Shiffrin, 1977), it is possible that this increases the attractiveness of the safer option. In support of this speculation, Charman and Howes (2003) showed that strategies that maximize consistency (such as copying and pasting a single item in computer use) are used even when more efficient (but less consistent) strategies become available. We hope to address this interpretation in a future study.

Note also that three of the problems (Problems 2, 8, and 10) involve choice among payoff distributions with two equally likely outcomes. Analysis of these problems can clarify the relationship of the current results to the results of animal timesaving decisions using similar distributions (Kacelnik & Bateson, 1996). The animal studies reflect risk seeking, which does not emerge here. The mean proportion of risky choices over the 50 trials of the three problems was 0.45 ($SD = 0.32$). The distance from random choice (the effect) was small ($d = 0.16$) and not significant, $t(59) = 1.17$, *ns*. Possible reasons for the inconsistency with animal studies are discussed below.

General Discussion

The current article highlights three properties of risk attitude in small timesaving decisions. The first two properties emerge when the attractiveness of the different outcomes can be reliably ranked. In these settings, small timesaving decisions exhibit the known properties of small money-related decisions (Erev & Barron, 2005): underweighting of rare events and the hot-stove (stickiness) effect. Underweighting of rare events is reflected by a preference for the alternative that leads to a better outcome (shorter delay) most of the time even when this alternative does not minimize the expected delay. The hot-stove effect is reflected by the observation that when feedback is limited to the obtained payoff, the tendency to select the safer alternative increases with experience. The third property emerges in environments in which the information available to decision makers does not allow for reliable ranking of the different delays. In these settings, behavior moves toward random choice.

Relationship Between Human and Nonhuman Timesaving Decisions

Comparison of the current results with those from studies of nonhuman timesaving decisions (see for example Bateson &

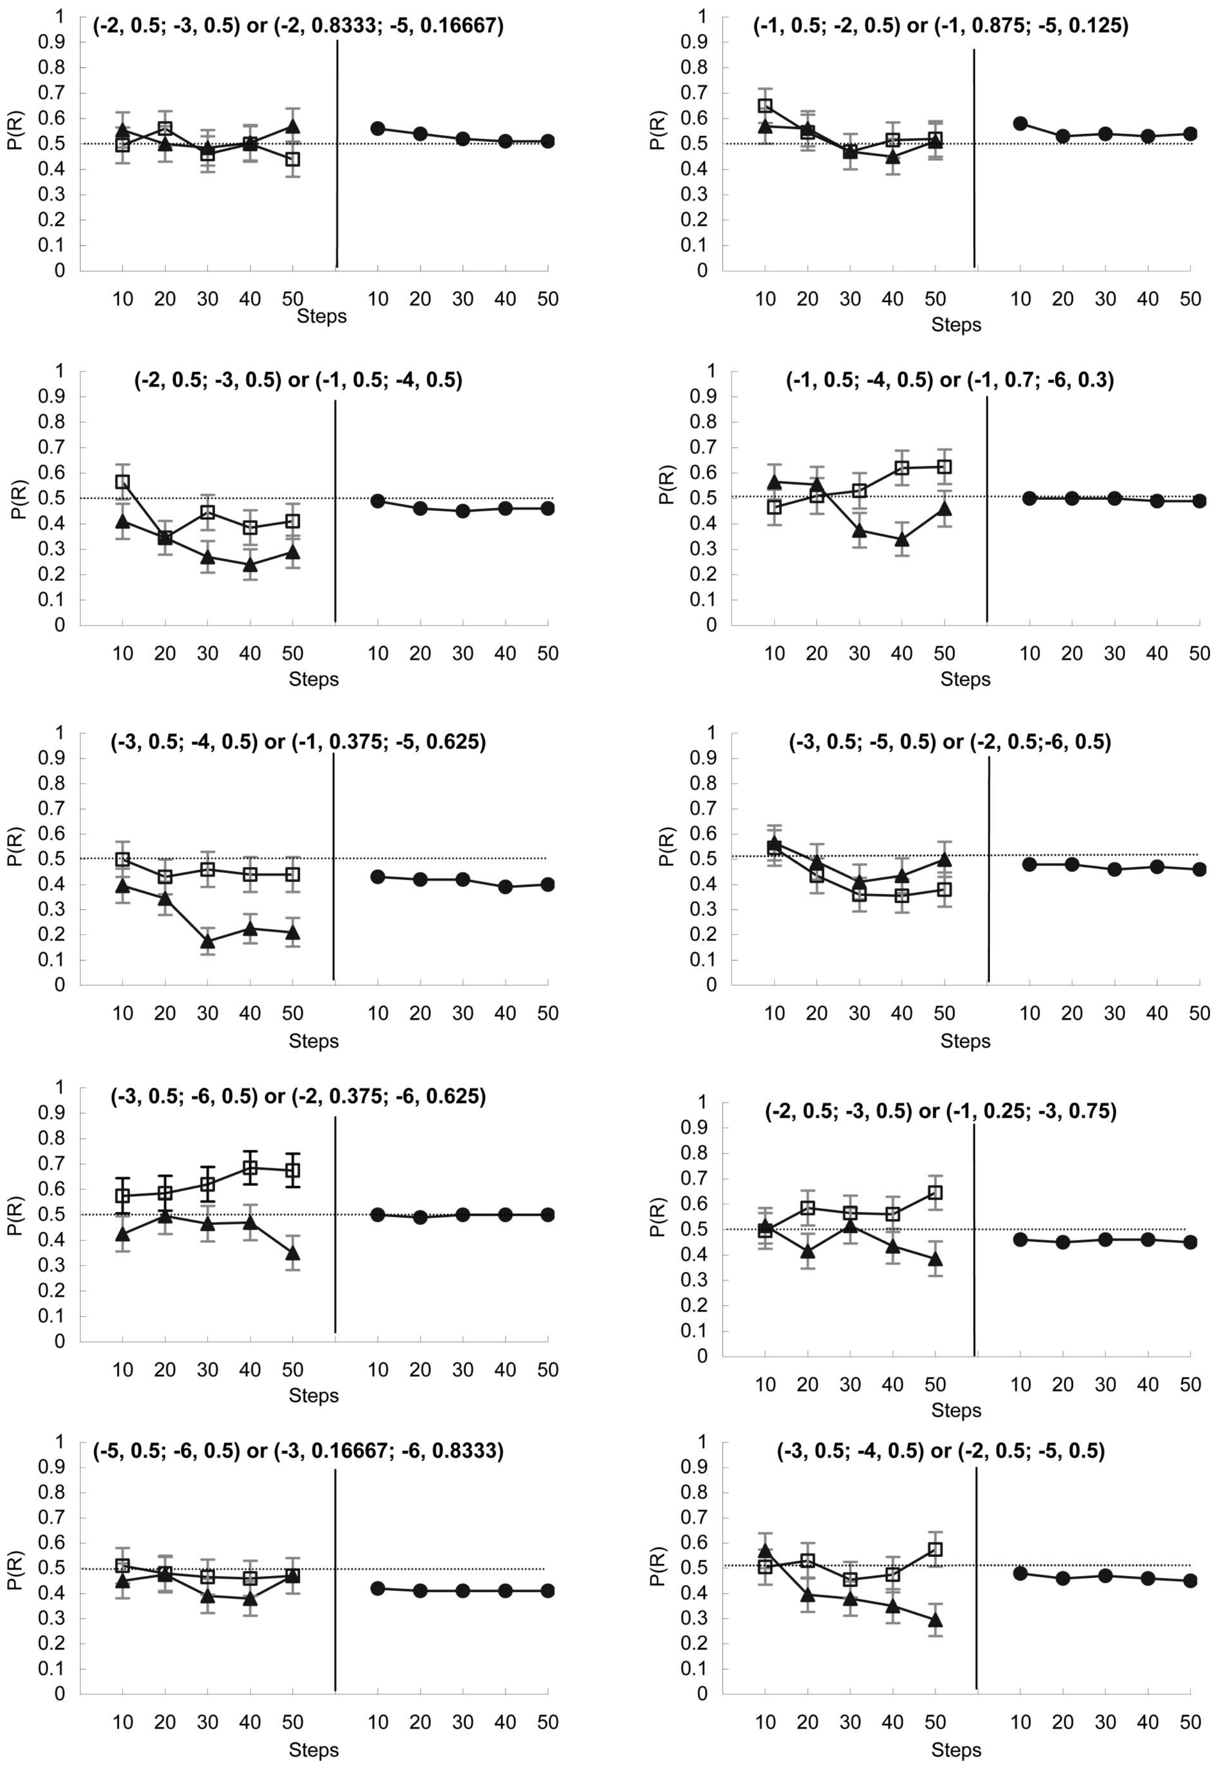
(-2, 0.5; -3, 0.5) or (-2, 0.8333; -5, 0.16667)
(-1, 0.5; -2, 0.5) or (-1, 0.875; -5, 0.125)
(-2, 0.5; -3, 0.5) or (-1, 0.5; -4, 0.5)
(-1, 0.5; -4, 0.5) or (-1, 0.7; -6, 0.3)
(-3, 0.5; -4, 0.5) or (-1, 0.375; -5, 0.625)
(-3, 0.5; -5, 0.5) or (-2, 0.5;-6, 0.5)
(-3, 0.5; -6, 0.5) or (-2, 0.375; -6, 0.625)
(-2, 0.5; -3, 0.5) or (-1, 0.25; -3, 0.75)
(-5, 0.5; -6, 0.5) or (-3, 0.16667; -6, 0.8333)
(-3, 0.5; -4, 0.5) or (-2, 0.5; -5, 0.5)
P(R)
0 0.1 0.2 0.3 0.4 0.5 0.6 0.7 0.8 0.9 1
10 20 30 40 50 10 20 30 40 50
Steps

Kacelnik, 1995; Marsh & Kacelnik, 2002) reveals an interesting difference. The literature suggests that animals exhibit risk seeking while trying to minimize delays. The general tendency for risk seeking was suggested to emerge in this case either as a by-product of associative learning (Kacelnik & Bateson, 1997) or by scalar utility theory (Reboreda & Kacelnik, 1991). Because studies of animal risk seeking in time delays are based on cases of symmetric variability (see Kacelnik & Bateson, 1997), the relevant comparison with our data should be with Problems 2, 8, and 10 of Experiment 4 (see Results of Experiment 4). In these problems, as in the animal studies, the variable outcomes were evenly distributed around the expected mean. Yet our human participants did not prefer the riskier alternative. One possible explanation for these different results is that our problems were not sufficiently similar to those applied in the animal studies. In the animal studies, the safer alternative was completely fixed (certain outcomes), whereas in our problems it was only less variable. However, theoretical explanations of animal data (the associative learning model and scalar utility theory) predict risk seeking also under these conditions (i.e., when one variable option is riskier than the other). The inconsistency with animal studies may therefore be related to other factors, perhaps the longer time delays applied in animal studies (in some cases, up to 60 s), or the different mechanisms applied by humans for ranking different delays. Counting steps, for example, was unlikely to be involved in animal studies but was probably applied by our human participants in Experiments 3 and 4. The mechanism of counting is not inconsistent with the assumptions of scalar utility theory about errors and memory representation or with those of the associative learning model (see Kacelnik & Bateson, 1997).

Relationship Between Decisions From Experience and Decisions From Description

Studies of scenario-based decisions (decisions from description) reveal interesting differences between money-related and timesaving decisions. Scenario-based money-related decisions tend to reflect risk seeking in the loss domain (see Kahneman & Tversky, 1979) when the decision concerns a single lottery. When the decision affects the outcome of multiple lotteries, choice behavior moves toward risk neutrality (see Keren & Wagenaar, 1987; Wedell & Bockenholt, 1994). Scenario-based studies of timesaving decisions tend to reflect risk aversion (e.g., Leclerc et al., 1995), especially when the outcomes affect other agents (see Krishnamurthy & Kumar, 2002).

The current analysis reveals that a different pattern emerges in the context of decisions from experience. In the current context the difference between money-related and timesaving decisions is not large. When the outcomes can be reliably ranked, both types of decisions reflect a tendency to prefer the option that leads to best outcomes most of the time and an increase in risk aversion with experience (when feedback is limited to the obtained payoff). It is important to emphasize, however, that the current results do not challenge the findings of previous scenario-based studies. The latter address decision tasks in which a complete description of the alternatives is available to the decision maker, whereas the current experimental study accounts for small feedback-based decisions. Notwithstanding, the differences between the results of the two types of studies suggest that the generalization of scenario-based studies to feedback-based timesaving decisions (and vice versa) should be considered with some caution.

Practical Implications and Future Research

The current analysis was motivated by the observation that many natural activities involve small timesaving decisions. We believe our results suggest that explicit study of these decisions can give rise to interesting implications. We conclude the article with two examples.

One example involves the optimal design of procedures aimed at saving data (creating backup files while editing a document, for example). Saving wastes time in most cases because the probability that the backup file will be needed is low. Thus, the current analysis suggests that people are likely to exhibit risk seeking in this context and to skip the saving operation. When saving is optimal (maximizes expected utility), this tendency implies counterproductive behavior. Under this logic, software developers, for example, could facilitate user performance by adding automatic saving (see supporting evidence in Yechiam, Haruvy, & Erev, 2002).

Another, more critical, example concerns traffic-related decisions. The current study offers a reasonable explanation for the tendency of drivers to run a red light, for example. Running a red light is a risky choice that generally leads to a good outcome for the actor: In most cases, the result will be minimized delays, whereas the potential costly consequences of this behavior—being stopped by the police or, worse, being involved in an accident—are rare. Because these rare outcomes can be so appalling, it is highly desirable to prevent drivers from preferring the option that leads to a better outcome most of the time. The current analysis suggests that the authorities responsible should prevent choices in this case from being based on learning from experience. One possibility is to expose drivers to scenarios of similar situations in which the choice leads to a negative outcome, on the grounds that when presented with scenarios, people demonstrate a tendency to make less risky choices (e.g. Leclerc et al., 1995; Krishnamurthy & Kumar, 2002)—even if they first experience the situation themselves (Barron, Leider, & Stack, 2006). Alternatively, if experience-based decisions cannot be avoided, our study suggests that increasing the frequency of negative experiences, such as getting a traffic ticket, may be key to changing behavior: The frequency with which violators are ticketed may be much more important than the cost of each fine.

Figure 4. Experiment 4: The left-hand column presents the proportion of risky choices in the 10 problems in the “actions” condition (—▲—) and in the “red light” condition (—⊟—). The right-hand column presents the predicted proportion of risky choices according to the reinforcement learning among cognitive strategies (RELACS) model (—◆—). Points represent the mean proportion over 10 trials; vertical lines depict 95% confidence intervals.

In summary, time is a valuable resource in modern life. Saving time, even if only short intervals are at stake, becomes an important factor in human decision making. We believe that exploring small time-related decisions may add to the understanding of key factors that underlie human behavior.

References

Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making, 16,* 215–233.

Barron, G., Leider, S., & Stack, J. (2006). *The effect of experience on a warnings impact* [working paper]. Boston: Harvard Business School.

Bateson, M., & Kacelnik, A. (1995). Preferences for fixed and variable food sources: Variability in amount and delay. *Journal of the Experimental Analysis of Behavior, 63,* 313–329.

Catalano, J. F. (1973). Effect of perceived proximity to end of task upon end-spurt. *Perceptual and Motor Skills, 36,* 363–372.

Chapman, G. B., Brewer, N. T., Coups, E. J., Brownlee, S., Leventhal, H., & Leventhal, E. A. (2001). Value for the future and preventive health behavior. *Journal of Experimental Psychology: Applied, 7,* 235–250.

Chapman, G. B., Nelson, R., & Hier, D. B. (1999). Familiarity and time preferences: Decision making about treatments for migraine headaches and Crohn's disease. *Journal of Experimental Psychology: Applied, 5,* 17–34.

Charman, S. C., & Howes, A. (2003). The adaptive user: An investigation into the cognitive and task constraints on the generation of new methods. *Journal of Experimental Psychology: Applied, 9,* 236–248.

Denrell, J., & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science, 12,* 523–538.

Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review, 112,* 912–931.

Estes, W. K. (1950). Toward a statistical theory of learning. *Psychological Review, 57,* 94–107.

Gibbon, J. (1991). Origins of scalar timing. *Learning and Motivation, 22,* 3–38.

Gray, W. D., & Boehm-Davis, D. A. (2000). Milliseconds matter: An introduction to microstrategies and to their use in describing and predicting interactive behavior. *Journal of Experimental Psychology: Applied, 6,* 322–335.

Grondin, S. (2001). From physical time to the first and second moments of psychological time. *Psychological Bulletin, 127,* 22–44.

Hornik, J., & Zakay, D. (1994). *Commercial time: The case of time and consumer behavior* [working paper]. Tel-Aviv, Israel: The Israel Institute of Business Research, Tel-Aviv University.

James, W. (1950). *The principles of psychology.* New York: Dover.

Kacelnik, A., & Bateson, M. (1996). Risky theories—The effects of variance on foraging decisions. *American Zoologist, 36,* 402–434.

Kacelnik, A., & Bateson, M. (1997). Risk sensitivity: Cross-roads for theories of decision-making. *Trends in Cognitive Sciences, 1,* 304–309.

Kacelnik, A., & Brito e Abreu, F. B. E. (1998). Risky choice and Weber's law. *Journal of Theoretical Biology, 194,* 289–298.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47,* 263–291.

Keren, G., & Wagenaar, W. A. (1987). Violation of utility theory in unique and repeated gambles. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13,* 387–391.

Kirby, K. N., & Herrnstein, R. J. (1995). Preference reversals due to myopic discounting of delayed rewards. *Psychological Science, 6,* 83–89.

Klatzky, R. L. (2000). When to inspect? Recurrent inspection decisions in a simulated risky environment. *Journal of Experimental Psychology: Applied, 6,* 222–235.

Krishnamurthy, P., & Kumar, P. (2002). Self–other discrepancies in waiting time decisions. *Organizational Behavior and Human Decision Processes, 87,* 207–226.

Leclerc, F., Schmitt, B. H., & Dubé, L. (1995). Waiting time and decision making: Is time like money? *Journal of Consumer Research, 22,* 110–119.

Lowenstein, G., & Elster, J. (1992). *Choice over time.* New York: Sage.

Lowenstein, G., & Thaler, R. H. (1989). Anomalities: Intertemporal choice. *Journal of Economic Perspectives, 3,* 181–193.

Marsh, B., & Kacelnik, A. (2002). Framing effects and risky decisions in starlings. *Proceedings of the National Academy of Sciences USA, 99,* 3352–3355.

Myers, J. L., & Sadler, E. (1960). Effects of range of payoffs as a variable in risk taking. *Journal of Experimental Psychology, 60,* 306–309.

Reboreda, J. C., & Kacelnik, A. (1991). Risk sensitivity in starlings: Variability in food amount and food delay. *Behavioral Ecology, 2,* 301–308.

Reyna, V. F., & Brainerd, C. J. (1995). Fuzzy-trace theory: An interim synthesis. *Learning and Individual Differences, 7,* 1–75.

Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review, 84,* 1–66.

Siegler, R. S., & Lemaire, P. (1997). Older and younger adults' strategy choices in multiplication: Testing predictions of ASCM using the choice/no-choice method. *Journal of Experimental Psychology: General, 126,* 71–92.

Tversky, A. (1969). Intransitivity of preferences. *Psychological Review, 76,* 31–48.

Weber, E. U., Shafir, S., & Blais, A. (2004). Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychological Review, 111,* 430–445.

Wedell, D. H., & Bockenholt, U. (1994). Contemplating single versus multiple encounters of a risky prospect. *American Journal of Psychology, 107,* 499–519.

Yechiam, E., Haruvy, E., & Erev, I. (2002). Reinforcement learning and the prevention of data catastrophes. *Journal of Managerial Psychology, 17,* 599–611.

Appendix A

Instructions for Participants: Experiment 1

In this experiment you will be operating a two-button money machine. On pressing a button, you will lose a sum of money which is a function of the time loss resulting from pressing the button. Your goal is to complete the experiment with as much money still remaining as possible.

Two squares will appear following each choice. Each square will last for a specific duration of time. The square on the *chosen* button (surrounded by a thin black frame) indicates the duration of time you *have lost*. The square on the *non-chosen* button (unframed) indicates the duration of time you *would have lost* if you had selected this button.

Note that as long as there is a square showing on one of the buttons you cannot press any button.

The basic payment is 300 Sheqels [approximately $67 U.S. dollars]. Your final payment is composed of the basic payment minus the sum you lose [1 second = 1 Sheqel].

Good luck.

Appendix B

The Experimental Screen: Experiment 1

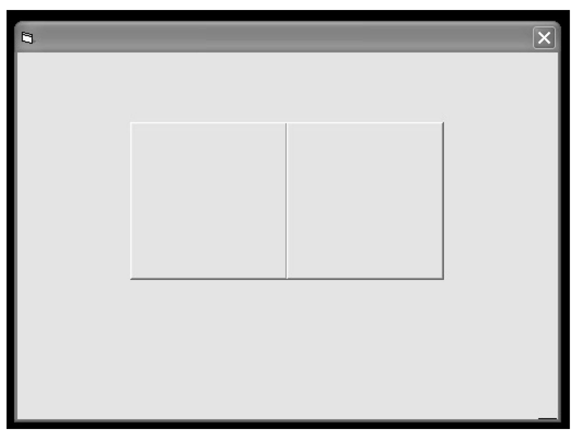

Figure B1. The experimental screen in Experiment 1. The participants were asked to select among the two unmarked buttons presented on a computer screen.

Appendix C

Instructions for Participants: Experiment 3

In this experiment you will be playing two independent games. In each of these games you have to choose one of two different paths by pressing one of the arrow keys on the keyboard (← for the left path, → for the right path). Then press the 'enter' key (↵).

Pressing the 'enter' key initiates exposure of the chosen path, composed of squares. Please press the 'enter' key as each square appears in order to expose the rest of the path until you reach the target point.

After you reach the target point, the words 'Choose a path' will be displayed at the top of the screen. You can then start the process again.

At the end of the first game, the words 'First Game Over' will be displayed on the screen, and the second game will automatically start. At the end of the experiment the words 'Game Over' will be displayed on the screen. You can leave the lab immediately after completing the experimental task. Thus, your goal is to complete the experiment as swiftly as possible.

Note that each square appears one second after you press the 'enter' key. Pressing either key during this second interval has no effect on the flow of the game.

Good luck.

(*Appendixes continue*)

Appendix D

The Experimental Screen: Experiment 3

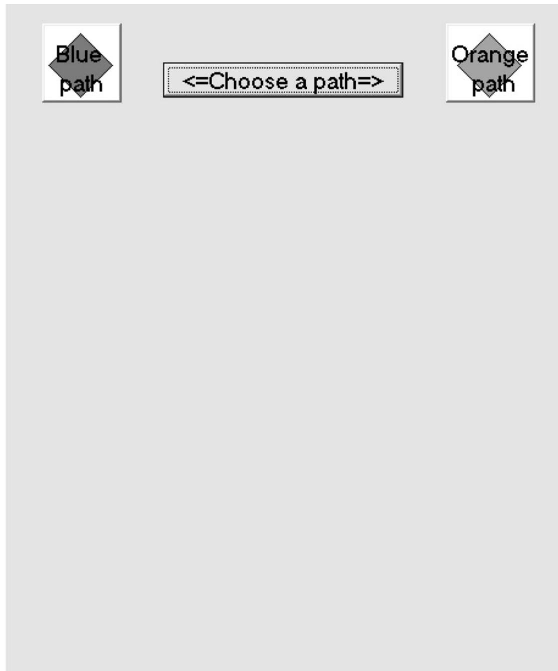

Figure D1. The experimental screen in Experiment 3. The participants were asked to select among the two paths.

Appendix E

Erev and Barron's (2005) Learning Model: Reinforcement Learning Among Cognitive Strategies (RELACS)

When all the possible outcomes have the same payoff sign (the current case), the model can be described by the following assumptions:

Assumption 1

In certain trials the decision maker follows a "fast best reply" strategy, which implies a selection of the action with the highest recent payoff. The "recent payoff" of action j is:

$$R_j(t + 1) = R_j(t)[1 - \beta] + v(t)_j \beta, \tag{1}$$

where $v(t)$ is the observed payoff from j in trial t, and β $(0 < \beta < 1)$ is a recency parameter: Large values imply large recency. Random choice is assumed when both actions have the same recent payoff. The recent value of action j is not updated in trials in which the payoff from j is not observed. The initial value $R_j(1)$ is assumed to equal the expected payoff from random choice.

Assumption 2

A second strategy considered by the decision maker is "case based." When this strategy is used before observing at least one outcome from each action (or if all previous outcomes were identical), it implies random choice. In other situations, one of the previous trials is selected and the action with the best payoff in that trial is selected (when forgone payoffs are not available to the agents, decisions are based on one random trial for each action). Ties are resolved on the basis of additional draws.

Assumption 3

The third strategy considered by the decision maker can be abstracted as a "slow best reply" rule. This strategy assumes a stochastic response rule that implies continuous but diminishing exploration. The probability that alternative (action) j is taken at trial t is:

$$p_j(t) = e^{W_j(t)\lambda/S(t)} / \sum_{k=1}^{2} (e^{W_k(t)\lambda/S(t)}), \tag{2}$$

where λ is an exploitation/exploration parameter (low values imply more exploration), $W_j(t)$ is the weighted average payoffs associated with alternative j, and $S(t)$ is a measure of payoff variability.

The weighted average rewards are computed like the recent payoffs, with the exception of a slower updating parameter α $(0 < \alpha < \beta)$, that is, $W_j(1) = R_j(1)$ and

$$W_j(t + 1) = W_j(t)[1 - \alpha] + v(t)_j\alpha. \tag{3}$$

The initial value of the payoff variability term, $S(1)$, is computed as the expected absolute difference between the obtained and expected payoff from random choice. The payoff variability measure moves toward the observed mean absolute difference between the obtained payoff $v(t)$ and the maximum of the last observed payoffs from the two actions ($Last_1$ and $Last_2$), that is,

$$S(t + 1) = S(t)[1 - \alpha] + ABS[v(t) - Max(Last_1, Last_2)]\alpha. \quad (4)$$

Assumption 4

Choice among strategies follows the stochastic choice rule described in Assumption 3, with one exception: The strategy's weighted average is updated only in trials in which the strategy was used. As in Assumption 3, the initial values are assumed to equal the expected payoff from random choice.

Parameters

The model has three parameters. Erev and Barron's (2005) estimation (based on the 40 money-related decisions they considered) yields the values $\lambda = 8$, $\alpha = 0.00125$, $\beta = 0.2$, and $\kappa = 4$.

Received May 19, 2005
Revision received May 17, 2006
Accepted May 31, 2006 ■