

Small Feedback-based Decisions and Their Limited Correspondence to Description-based Decisions

GREG BARRON and IDO EREV*

Faculty of Industrial Engineering and Management, Technion, Haifa, Israel

ABSTRACT

The present paper explores situations in which the information available to decision makers is limited to feedback concerning the outcomes of their previous decisions. The results reveal that experience in these situations can lead to deviations from maximization in the opposite direction of the deviations observed when the decisions are made based on a description of the choice problem. Experience was found to lead to a reversed common ratio/certainty effect, more risk seeking in the gain than in the loss domain, and to an underweighting of small probabilities. Only one of the examined properties of description-based decisions, loss aversion, seems to emerge robustly in these ‘feedback-based’ decisions. These results are summarized with a simple model that illustrates that all the unique properties of feedback-based decisions can be a product of a tendency to rely on recent outcomes. Copyright © 2003 John Wiley & Sons, Ltd.

KEY WORDS Probability learning; feedback-based decisions; reinforcement learning; prospect theory; Allais paradox

Many common activities involve ‘small’ decision problems. Driving, for example, requires repeated selection among routes, speeds, and various other options. Although little time and effort is typically invested in these and similar small decisions, they can be consequential. The estimated cost of traffic accidents in the USA is more than 100 billion dollars a year (see e.g. Blincoe, 1994¹), and many of the accidents are at least partially products of *ex-post* unwise decisions.

The current paper focuses on an important subset of the small decision problems exemplified above that can be referred to as ‘small feedback-based’ decisions. These problems are defined by three main properties. First, they are repeated; decision makers face the same problem many times in similar situations. Second,

* Correspondence to: Ido Erev, Faculty of Industrial Engineering and Management, Technion, Haifa, Israel.
E-mail: erev@tx.technion.ac.il

Contract/grant sponsors: USA–Israel Binational Science Foundation; Technion VPR; Max Wertheimer Minerva Center.

¹The 1994 estimate is 150.5 billion. This estimate represents the present value of lifetime economic costs for 40,676 fatalities, 5.2 million non-fatal injuries and 27 million damaged vehicles.

each single choice is not very important; the alternatives tend to have similar expected values that may be quite small. Finally, the decision makers do not have 'objective' prior information concerning the payoff distributions. In selecting among the possible options, they have to rely on the immediate feedback obtained in similar situations in the past.²

Notice that small feedback-based decision problems are quite distinct from the problems studied in mainstream decision research. Mainstream research tends to focus on 'big' decisions that are made based on a careful evaluation of the possible payoff distributions. These decisions are naturally studied in a 'description-based' paradigm. In a typical study (e.g. the studies used by Kahneman & Tversky, 1979, to validate prospect theory) the participants are presented with a complete description of a non-trivial choice problem and are asked to make one selection among the possible options.

One reasonable approach to modeling feedback-based decisions assumes that decision makers use the obtained feedback to discover the underlying choice problem (to estimate the payoff distributions of the different alternatives). Under this assumption, the large structural differences between feedback-based and description-based decision problems may not lead to large behavioral differences. In both cases the decision makers are expected to respond to the underlying payoff distributions. This interesting assumption was supported in research conducted by Benartzi and Thaler (1995). They show that prospect theory, the popular and elegant summary of the main results obtained in studies of description-based decisions, can account for real-life feedback-based investment decisions. In particular, it explains the suggestion that most people are under-invested in the stock market relative to the market's past performance. According to this account, loss aversion, one of the main concepts in prospect theory, leads investors to prefer safe prospects that eliminate losses even at a very high cost (in terms of expected return). This explanation was later supported by the findings of a laboratory experiment (Thaler et al., 1997; and see Gneezy & Potters, 1997) that shows the pattern predicted by prospect theory emerging in 200 repetitions of a controlled investment task.³

However, there are also good reasons to predict important differences between feedback-based and description-based decisions. The most obvious reason involves the assertion that immediate feedback is likely to lead to adaptive learning that moves behavior toward expected value maximization. This assertion is consistent with certain interpretations of the law of effect (Thorndike, 1898) and with the view that 'practice makes perfect'. Since studies of description-based decisions reveal robust deviations from maximization, the 'adaptive learning' assertion implies higher maximization rates in feedback-based decisions.

The hypothesis of higher maximization rates in feedback-based decisions can also be supported based on previous studies of decisions among repeated play of gambles (see Lopes, 1981; Keren & Wagenaar, 1987; Keren, 1991; Bockenholt & Wedell, 1990). In a typical study, decision makers were asked to select which of two gambles they prefer to play 10 times. Comparison of the results with a control condition in which the selected gamble was played only once showed that the expectation of playing the gambles repeatedly increased the tendency to select the gamble with the highest expected value.⁴ Although feedback-based decisions differ from these 'repeated-play' conditions in many respects, they involve repeated play. Thus, a similar trend can be predicted.

A very different type of difference is predicted under a variant of the assertion, made above, that decision makers simply use feedback to estimate the payoff distributions. If the samples used to derive the estimates are relatively small, the estimates will not be accurate. Inaccurate estimates that can lead to large differences between feedback- and description-based decisions are particularly likely when rare events are important and the payoff distributions have large variances.

²The availability of immediate and unbiased feedback distinguishes the problems studied here from the problems used by Herrnstein and his colleagues (see e.g. Herrnstein et al., 1993) to study melioration.

³Although real-life investment decisions may not be a good example of small decisions, Thaler et al.'s experiment did study small decisions. The maximal expected payoff per trial was 'only a few cents' (Thaler et al., 1997).

⁴Tversky and Bar Hillel (1983) note that this effect can be a result of the change in the implied payoff distributions.

A fourth reason to predict unique properties of feedback-based decisions comes from previous studies of decision making in signal detection tasks (see review in Erev, 1998). These studies show regularities that seem very different from the regularities observed in description-based decisions. For example, Barkan et al. (1998) documented robust underweighting of rare events (that tend to be overweighted in description-based studies). Since signal detection tasks are similar to feedback-based tasks (with the exception of the availability of trial specific information), it is natural to assume that this pattern will be observed in feedback-based decisions.

Finally, the study of choice behavior in repeated games reveals high sensitivity to recent feedback (see e.g. Erev & Roth, 1998; Cheung & Friedman, 1999; Camerer & Ho, 1999; Selten & Buchta, 1998; Daniel et al., 1998; Cooper et al., 1997; Chen & Tang, 1999). This observation is inconsistent with the assumption that the feedback is effectively used to estimate the relevant distributions.

The main goal of the current research is to improve our understanding of the differences between one-shot description-based and small feedback-based decisions. To achieve this goal we chose to focus on four of the best-known properties of choice behavior in one-shot description-based tasks and examined if similar properties emerge when decision makers have to base their choices on the obtained feedback.

In an attempt to study simple small feedback-based decisions, the current study examines binary choice tasks where the only relevant information available is the obtained payoff from previous decisions. In each trial of the current studies (with the exception of Experiment 5) the decision makers were asked to choose between the two keys of 'a computerized money machine' (see the Appendix). Each selection yielded a payoff in points—a draw from the selected key's payoff distribution. This payoff and the accumulated earnings were presented on the screen and constituted the only information available to the decision makers. The decision makers were told that their goal was to maximize their number of points and that these points would determine their final payoff. They neither received prior information concerning the payoff distributions, nor did they know that the payoffs were random draws from static payoff distributions.⁵

This paper is organized as follows. Experiment 1 presents a replication of Thaler et al. (1997). It shows that the loss aversion tendency they documented can be observed in the minimalist paradigm used here.

Experiments 2, 3, and 4 show that, unlike the loss aversion phenomenon, other properties of one-shot description-based decisions do not necessarily emerge when decision makers rely on feedback. Interestingly, the exact opposite tendencies were observed. Experiment 2 reveals that experience with feedback-based decisions can make an outcome provided with certainty less attractive. This trend is a mirror image of the common ratio/certainty effect captured by prospect theory. Experiment 3 shows that feedback-based decisions reflect a stronger risk aversion tendency in the loss domain than in the gain domain. This pattern is a reversal of the pattern observed in description-based studies. Experiment 4 demonstrates that when decision makers rely on feedback, they tend to underweight rare outcomes. This tendency is the opposite of the overweighting of rare outcomes observed in description-based studies.

The relationship of these results to the decision-making literature is clarified in Experiment 5 and with a simple model. Experiment 5 suggests that the effect of feedback is quite distinct from the effect of the expectation of playing the selected gamble repeatedly. The model demonstrates that the unique properties of feedback-based decisions can be a product of a tendency to rely on recent outcomes.

EXPERIMENT 1: LOSS AVERSION (A REPLICATION OF THALER ET AL., 1997)

A comparison of two of the conditions examined by Thaler et al. (1997) provides evidence for loss aversion in a feedback-based decision task. In Condition 1 ('Monthly' in their scenario) subjects repeatedly divided a

⁵The main cost of this minimal information setting involves the difficulty in deriving 'rational behavior' (any behavior can be rational under certain prior beliefs). Since the current research is not designed to ask if people are 'rational', we are not concerned by this cost (see Budescu et al., 1997, for a discussion of the value of 'rationality free' decision research).

portfolio between a risky and a safe investment. The risky investment ('stocks') produced high-yield returns with large variability (a normal payoff distribution with a mean of 1 and standard deviation of 3.54). The safe investment ('bonds') produced low-yield returns but insured small variability and positive outcomes (a normal distribution with a mean of 0.25 and standard deviation of 0.177 that was truncated at 0). Subjects were not informed as to the actual payoff matrix and were told only that they were dividing their portfolio between two unnamed funds. In Condition 2 ('Inflated monthly') payoffs from both investments were increased by a constant so that all payoffs were positive. The results reveal that in Condition 1 subjects consistently preferred the low-risk (and low expected value) bonds. Interestingly, 200 trials with immediate feedback did not lead subjects to prefer the stocks and to maximize expected return. Over the 200 trials, the proportion of expected-value maximizing choices was 0.42. When returns were all in the gain domain (Condition 2), the maximization rate increased to 0.56. Thaler et al. note that these results are consistent with the prediction of a natural generalization of prospect theory to repeated decision tasks.

To derive the implications of these results for small feedback-based decisions it is important to distinguish between two possible interpretations. Under the first, the observed loss aversion pattern is facilitated by the investment scenario and is robust to the availability of feedback. Under a second interpretation, loss aversion emerges from feedback. Experiment 1 was designed to evaluate the second interpretation. It is a simplified replication of the conditions described above, with the addition of a condition in which the variance of both distributions was small. Subjects partook in a repeated binary task where they chose between two buttons without receiving prior information as to the payoff distribution of either button. This minimalist design was implemented to evaluate the pure effect of feedback.

Design

Three payoff conditions were compared. The basic task in all three conditions was a binary choice between two approximately normal payoff distributions referred to as L and H. This basic task was performed 200 times (with immediate feedback) in each of the conditions. The conditions differed with respect to the two distributions. Following Thaler et al., in the baseline Condition (Problem 1), distribution L was centered at 25, had a standard deviation of 17.7, and was truncated at 0 to prevent negative returns.⁶ The implied mean was 25.63. Distribution H had a mean of 100 points with a standard deviation of 354. In the second condition (Problem 2), all payoffs were transformed to the gain domain by adding 1200 points to the means of both distributions while the standard deviations remained unchanged. The third condition (Problem 3) was identical to the second except that the standard deviation of both distributions equaled 17.7. At the conclusion of the experiment, points were converted to monetary payoffs according to the exchange rate: 100 points = 0.05 Shekels (about 1.25 US cents).⁷

Subjects

Thirty-six Technion students served as paid participants in the experiment. Most of the participants in this and the other experiments described in this paper were second- and third-year industrial engineering and economics majors who had taken at least one probability and economics course. They were randomly assigned to one of the three experimental conditions. In addition to the performance contingent payoff, described above, subjects received 5 Shekels for showing up. Final payoffs ranged between 6.25 and 15 Shekels (\$2–\$4).

⁶The draws were rounded to the nearest integer. The other payoff distributions were not truncated.

⁷In Problems 2 and 3 subjects were only paid for points above 240,000 (1200 points \times 200 trials).

Apparatus and procedure

Subjects were informed that they were playing on a ‘computerized money machine’ (see a translation of the instructions in the Appendix) but received no prior information as to the game’s payoff structure. Their task was to select one of the ‘machines’ two unmarked buttons (see the figure in the Appendix) in each of the 200 trials. Participants were aware of the expected length of the experiment (30 minutes to an hour), so they knew that it included many rounds. To avoid an ‘end of task’ effect (e.g. a change in risk attitude), they were not informed that the experiment included exactly 200 trials.⁸ Payoffs were contingent upon the button chosen; they were drawn from the distribution associated with the selected button, described above. Two types of feedback immediately followed each choice: (1) the payoff for the choice, that appeared on the selected key for the duration of 1 second, and (2) an update of an accumulating payoff counter, which was constantly displayed.

Results

The graph in the left-hand cell of Figure 1 presents the proportion of maximization (H) choices in 20 blocks of 10 trials over the 12 participants in each condition (the right-hand cell presents the predictions of the model discussed below). The learning curves show non-monotonic trends in Problems 1 and 2. The proportion of H choices decreased initially, and then increased. We return to this pattern below.

The aggregated results in Problems 1 and 2 are practically identical to those of the corresponding conditions in Thaler et al. In line with the counter-intuitive prediction made by a generalization of prospect theory, the average subject in Problem 1 did not learn to prefer the alternative that maximizes expected value. Over the 200 trials, subjects’ aggregate proportion of H choices was 0.30 (STD = 0.22) in Problem 1, 0.51 (STD = 0.21) in Problem 2, and 0.85 (STD = 0.08) in Problem 3. The differences are significant ($t[22] = 2.43$, $p < 0.05$ for Problems 1 and 2, and $t[22] = 5.28$, $p < 0.001$, for Problems 2 and 3).

Notice that the difference between Problems 2 and 3 is an example of the payoff variability effect (see Myers et al., 1965; Busemeyer & Townsend, 1993): an increase in payoff variability impairs maximization.

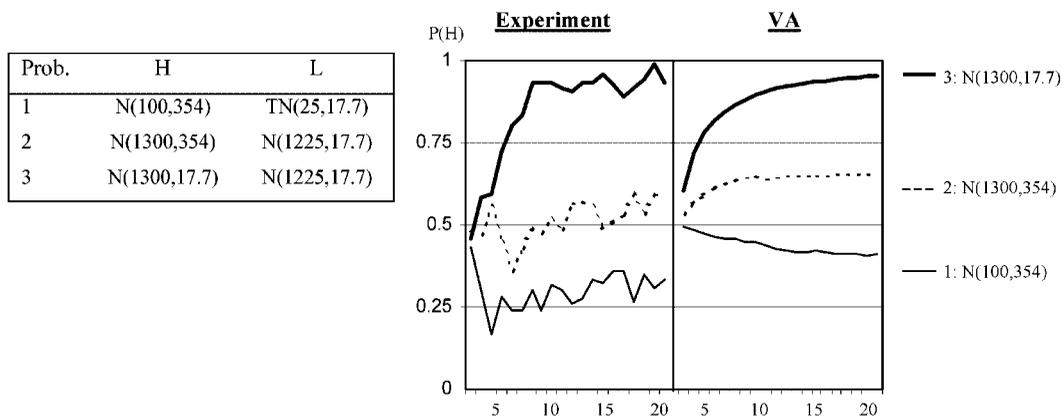


Figure 1. Replication of Thaler et al. (Experiment 1). Each curve in the left-hand column shows the observed proportion of maximization (H) choices in 20 blocks of 10 trials in one of the three problems. The problems are presented on the left: Alternative N(x , y) yields a draw from a normal distribution with mean of y and standard deviation x . In TN(x , y) the payoff is truncated at 0 (0 is the worst possible payoff). The right-hand column presents the predictions of the VA model

⁸Not knowing the length of the experiment also prevents subjects from using probability-based reasoning (the focus on the likelihood of achieving a particular aspiration level) (Lopes, 1996). This type of reasoning bases choice on the probability of coming out ahead, which is a function of the number of choices to be made. A second reason for not telling subjects the game’s length is that this better approximates the real-world small decisions that interest us. In such situations, the number of future choices to be made is often unknown.

This difference suggests that there are two reasons for the low maximization rate in Problem 1: loss aversion and high payoff variability.

The observation that an increase in payoff variability impairs maximization is, of course, not surprising; even optimal procedures designed to find the maximum are less accurate when the variability increases. Nevertheless, the behavioral payoff variability effect has many interesting implications. One set of implications is concerned with the effect's magnitude and its robustness to experience. As Figure 1 shows, the large effect does not seem to decrease with time. Experiment 2 evaluates another nontrivial implication.

EXPERIMENT 2: THE REVERSED COMMON RATIO/CERTAINTY EFFECT

Allais (1953) has noticed that the more risky of two prospects becomes relatively more attractive when the probability of winning in both prospects is multiplied by a common ratio. This 'common ratio' effect is important because it represents a robust violation of expected utility theory's substitution principle. Kahneman and Tversky (1979) demonstrate this effect with an experimental study of the following choice problems (the outcomes represented hypothetical payoffs in thousand Israeli Lira):

Problem 4. Choose between:

L: 3 with certainty

H: 4 with probability 0.8; 0 otherwise

Problem 5: Choose between:

L: 3 with probability 0.25; 0 otherwise

H: 4 with probability 0.2; 0 otherwise

Notice that Problem 5 was created by dividing the probability of winning in Problem 4 by four. Kahneman and Tversky's results revealed that while 80% of the subjects preferred L in Problem 4, only 35% preferred L in Problem 5. In prospect theory the common ratio effect is an example of the certainty effect, and it is captured by the assumption of a nonlinear probability weighting function that overweighs outcomes that occur with small probabilities (and underweighs outcomes that occur with high probabilities). Specifically, Alternative H is unattractive in Problem 4 because the '0' outcome ($p = 0.2$) is overweighted, but it is attractive in Problem 5 because the '4' outcome ($p = 0.2$) is overweighted. Interestingly, MacDonald et al. (1991) found a similar pattern in feedback-based decisions by rats.

The current experimental paradigm appears to fall between the paradigms studied by Kahneman and Tversky and MacDonald et al. We focus on human subjects (like Kahneman & Tversky) and on feedback-based decisions (like MacDonald et al.). Since Kahneman and Tversky and MacDonald et al. found the same pattern, interpolation predicts that a similar pattern will be found here. However, a different prediction can be derived from the payoff variability effect. Payoff variability (relative to expected payoff difference) is larger in Problem 4 than in Problem 5. Thus, the robust payoff variability effect found in Experiment 1 predicts a higher maximization rate in Problem 4 than in Problem 5. The current study compares these competing predictions.

Design

Problems 4 and 5 were compared. The translation from points to monetary payoffs was according to the exchange rate: 1 point = 0.01 Shekel (0.25 US cents).

Subjects

Forty-eight Technion students served as paid participants in the experiment. They were randomly assigned to one of the two experimental conditions. Subjects received an initial (showing-up) fee and additional payoff

contingent upon performance. The total payoff ranged between 20 and 30 Shekels (\$5–7). The showing-up fee was 12.60 Shekels in Problem 4 and 21.90 Shekels in Problem 5; the difference was implemented to equate the expected earnings in both conditions assuming random choice.

Apparatus and procedure

The procedure was identical to Experiment 1 with the exception that the experiment was run for 400 trials.

Results

The experimental data are presented in Figure 2. The mean proportion of H choices (high EV but more risky alternative), over subjects and blocks, was 0.63 (STD = 0.24) when L provided ‘3 with certainty’, it decreased significantly to 0.51 (STD = 0.10) when payoff probabilities were divided by four ($t[46] = -2.35, p < 0.05$). Thus, as predicted by the payoff variability hypothesis, the aggregated choice proportion deviated from the predictions of identical choice proportion in the two problems in the opposite direction of the deviation observed in description-based decisions. We refer to this deviation as the ‘reversed common ratio/certainty’ effect.

Notice that a comparison of the current results to the results reported by MacDonald et al. appear to suggest that there are differences between humans and rats. However, the difference in subject population is only one of many differences between the two experiments. For example, MacDonald et al. used a within-subject design and allowed the decision makers to immediately consume their rewards.

EXPERIMENTS 3a AND 3b: THE REVERSED PAYOFF DOMAIN EFFECT

Experimental studies of decision making in one-shot description-based tasks reveal a ‘payoff domain effect’: Decision makers tend to be more risk averse in the gain domain than in the loss domain. A particularly strong example, also known as the ‘reflection effect’, was documented by Kahneman and Tversky (1979) in a comparison of Problem 4 (see above), with:

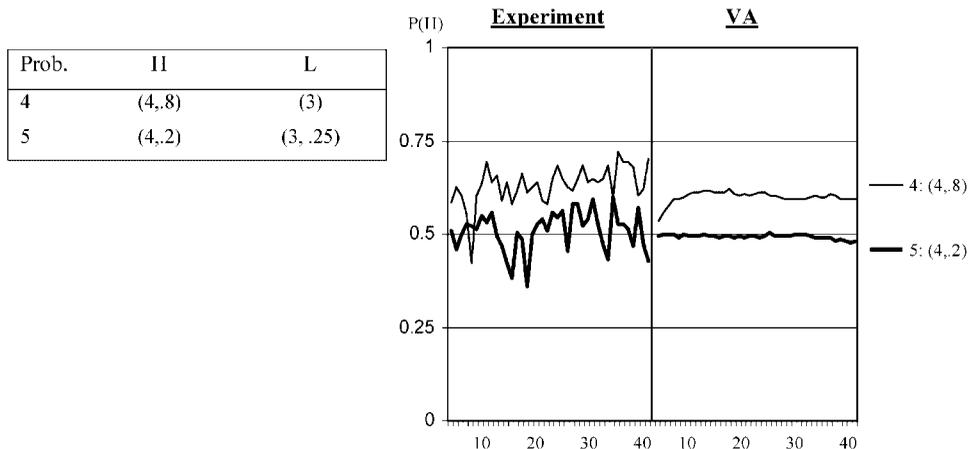


Figure 2. The reversed certainty effect (Experiment 2). Each curve in the left-hand column shows the observed proportion of maximization (H) choices in 40 blocks of 10 trials in one of the two problems. The problems are described on the left: Alternative (V, p) yields payoff of V with probability p, and zero otherwise. Alternative (V) yields payoff of V with certainty. The right-hand column presents the predictions of the VA model

Problem 6. Choose between:

L: -4 with probability 0.8; 0 otherwise

H: -3 with certainty

Kahneman and Tversky's results revealed deviation from expected value maximization (preference of L) in both problems: 80% of the subjects preferred L in Problem 4, and 92% preferred L in Problem 6. The deviation from maximization in the gain domain (Problem 4) can be described as risk aversion, and the deviation in the loss domain (Problem 6) can be described as risk seeking. Prospect theory captures both deviations with the assumption of an S-shaped value function. The present experiments examine if a similar pattern emerges in feedback-based decisions.

Experiment 3a

Design

Two choice problems were compared: Problem 4 from Experiment 2, and Problem 6 (note that new subjects partook only in Problem 6). As in Experiment 2, the exchange rate was 1 point = 0.01 Shekel (0.25 US cents). To insure similar objective incentive structures in the two conditions, the initial showing-up fee in Problem 6 was 37.4 Shekels.

Subjects, apparatus, and procedure

As in Experiment 2.

Results

The experimental results for the two conditions, presented in the top panel of Figure 3, show maximization rates of more than 0.5. Thus, feedback did not lead to more risky choices in the loss than the gain domain. In fact, the opposite pattern was observed. The proportion of risky choices was 0.63 (STD = 0.24) when prospects were positive, and it decreased significantly to 0.40 (STD = 0.19) when prospects were negative (L in Problem 4) ($t[46] = 3.64, p < 0.05$). The difference between the maximization rates is insignificant ($t[46] = 0.58, NS$).

Experiment 3b

Although the results of Experiment 3a are a mirror image of the famous demonstration of the payoff domain effect, they do not prove that feedback can reverse this effect. Recall that the low rate of H choices in the description-based decisions in Problems 4 and 6 implies a deviation from maximization in the direction of the payoff domain effect. The higher rate of H choices observed in feedback-based decisions can, therefore, be interpreted in two ways: a stronger tendency to maximize earning, and/or a reversed payoff domain effect. To evaluate if a clear reversed payoff domain effect can emerge, we chose to study choice among gambles with equal expected values.

Design

Two binary choice problems were compared (see lower panel in Figure 3). Problem 7 involves a choice between a safe prospect (S) that promises '9 with certainty' and a risky prospect (R) that yields '10 with $p = 0.9$; 0 otherwise'. In Problem 8 the safer prospect (S) yields '-9 with certainty' and the risky prospect (R) yields '-10 with $p = 0.9$; 0 otherwise'. As in Experiment 2, the exchange rate was 1 point = 0.01 Shekel. To insure similar objective incentive structures in the two conditions, the initial showing-up fee in was set at 7 Shekels in Problem 7 and at 43 Shekels in Problem 8.

Subjects, apparatus, and procedure

As in Experiment 2 with the exception that 30 subjects participated in each condition and the experiment lasted only 200 trials.

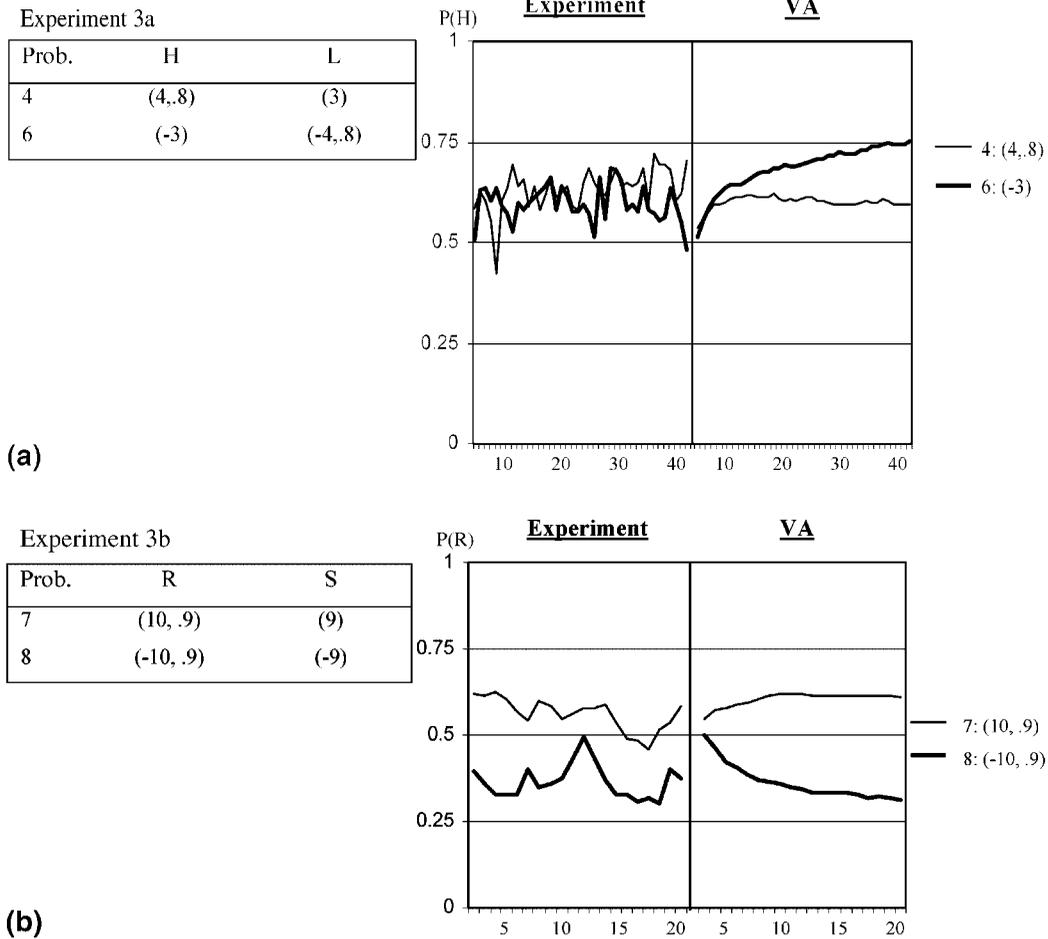


Figure 3. The reversed payoff domain effect (Experiments 3a and 3b). (a) Observed (left) and predicted (right) proportion of maximization (H) choices in 40 blocks of 10 trials. (b) Observed (left) and predicted (right) proportion of risky (R) choices in 20 blocks of 10 trials

Results

The results are presented in the lower panel of Figure 3. Over trials and subjects the proportion of risky choice (preferring the gamble over the sure payoff) was 0.56 (STD = 0.30) in the gain domain and 0.37 (STD = 0.27) in the loss domain. The difference between the two proportions is significant ($t[58] = 2.65, p < 0.01$). Thus, in the current setting, decision makers appear to take more risk in the gain than in the loss domain. We refer to this pattern as the ‘reversed payoff domain effect’.

EXPERIMENT 4: UNDERWEIGHTING OF SMALL PROBABILITIES

In one-shot description-based tasks decision makers behave ‘as if’ they overweight small probabilities. For example, Kahneman and Tversky (1979) found that most people prefer the gamble (5000, with $p = 0.001$; 0 otherwise) over a sure payoff with the same expected value. These findings can be captured by with the assumption of a nonlinear weighting function that implies the overweighting of small probabilities

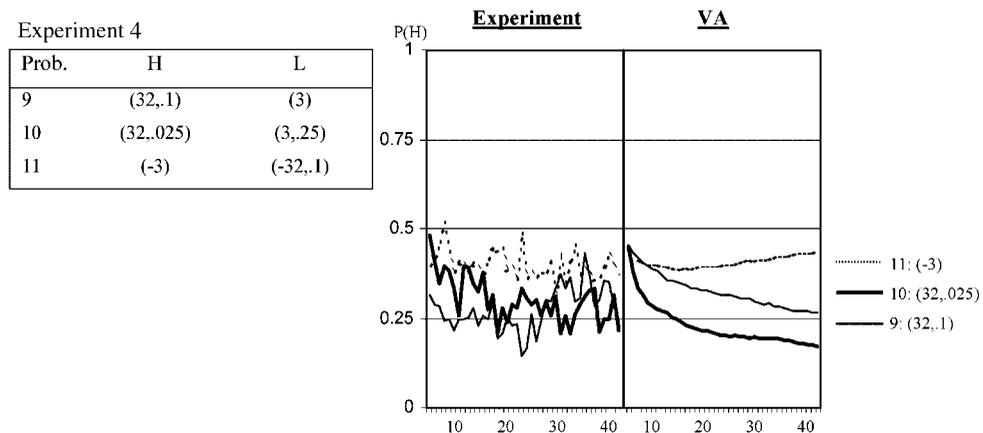


Figure 4. Underweighting of small probabilities (Experiment 4). Observed (left) and predicted (right) proportion of maximization (H) choices in 40 blocks of 10 trials

(in prospect theory), and with the idea that the rare events ‘tax’ the more likely events (in the TAX model, see Birnbaum & Navarrete, 1998). Careful quantification of this effect shows that the average decision maker behaves as if she overweighs probabilities below 0.25 (see a recent quantification of this effect in Gonzalez & Wu, 1999). To examine whether feedback-based decisions lead to a similar pattern, we examined 400 trials of play in three problems containing small probabilities.

Design

The three choice problems presented in Figure 4 were examined. As in Experiments 2 and 3, the exchange rate was 1 point = 0.01 Shekel. The initial showing-up fee was 12.60 Shekels in Problem 9, 21.90 Shekels in Problem 10, and 37.40 Shekels in Problem 11.

Subjects, apparatus, and procedure

As in Experiment 2.

Results

Figure 4 shows the observed learning curves. Over the 400 repetitions the proportions of maximization (H) choices were 0.28 (STD = 0.28), 0.30 (STD = 0.21), and 0.40 (STD = 0.25) in Problems 9, 10 and 11 respectively. These proportions are not different from each other ($F(2, 46) = 1.63$, ns), and the grand mean is significantly below the 0.50 maximization rate predicted under random behavior ($t[71] = 5.88$, $p < 0.05$).

These results are particularly interesting because, in the current problems, the overweighting of rare events observed in description-based decisions would imply maximization. Yet our results show the opposite trend in both the gain and loss domains that can be described (using the terminology of prospect theory) as the emergence of underweighting of rare events.

Another interesting implication of the current results is concerned with the effect of the expectation of playing the gambles repeatedly. Recall that previous research shows in description-based decisions this expectation increases maximization. The low maximization rates observed here suggest that experience with repeated choice (in feedback-based decisions) can have a different effect.

EXPERIMENT 5: DESCRIPTION-BASED CHOICE

In Experiments 1–4 we compared the known regularities observed in previous description-based studies to the results observed in the new feedback-based conditions. This convention was adopted under the assumption that there is no reason to replicate the properties of description-based decisions (that have been documented in dozens of studies). Whereas this assumption seems reasonable, it might not be exactly correct. It is possible that at least some of the current findings reflect the unique properties of our subject population and not the unique properties of feedback-based decisions. To evaluate this hypothesis, Experiment 5 studies two versions of a description-based decision task. The ‘one-play’ condition was a replication of the studies used to validate prospect theory. The participants were presented with descriptions of pairs of hypothetical prospects, and were asked to select which prospect they prefer to play once in each pair. In condition ‘repeated play’ the participants were presented with the same stimuli, and were asked to select the prospect that would be played 100 times in each pair.

Design

Five of the problems examined above (Problems 4, 5, 6, 9, and 10) were compared.

Subjects, apparatus, and procedure

The participants were 161 Technion students. They were asked to complete a short questionnaire during class. Each participant was presented with all five problems under one of the two conditions. The ninety-one participants who were assigned to the one-play condition were instructed to imagine that each chosen gamble would be played once. Seventy participants were assigned to repeated-play and were asked to imagine that each selected gamble would be played 100 times.

Results

The ‘description-based’ columns of Table 1 present the proportion of maximization choices in the two conditions. The right-hand column compares these proportions to the maximization rates in the last 100 trials of the relevant condition in Experiments 2–4.

Notice that a comparison of Problem 4 and Problem 5 in the one-play condition shows the common ratio/certainty pattern. Gamble L is more attractive when the safe outcome is provided with certainty ($Z = 3.81$, $p < 0.05$, test for difference in proportions). Yet, the magnitude of the effect observed here is smaller than the effect observed by Kahneman and Tversky with larger amounts. Comparison of Problems 4 and 6 in the

Table 1. Proportions of maximization (H) choices in the five problems studied in Experiment 5

Problem	The two prospects		Description-based		Feedback-based (Last 100) $n = 24$
	L	H	One-play $n = 91$	Repeated-play $n = 70$	
4	(3)	(4, 0.8)	0.55	0.61	0.66
5	(3, 0.25)	(4, 0.2)	0.81 ^r	0.47	0.50
6	(−4, 0.8)	(−3)	0.30 ^f	0.49	0.58
9	(3)	(32, 0.1)	0.47	0.57 ^f	0.32
10	(3, 0.25)	(32, 0.025)	0.47	0.69 ^f	0.27

r—significantly different ($p < 0.05$) from Repeated-play.

f—significantly different ($p < 0.05$) from Feedback-based.

Note: The description-based columns show the results in the two questionnaire conditions. The right-hand column presents the proportion of maximization choices in the last 100 trials of the relevant conditions in Experiments 2–4.

one-play condition demonstrates the payoff domain effect. Participants appear to be risk seekers in the loss domain (Problem 6) but less so in the gain domain (Problem 4) ($Z = 3.45$, $p < 0.05$).

Comparison of the one-play to the repeated-play condition shows a slightly higher maximization rate in repeated-play (0.57 (STD = 0.29) versus 0.52 (STD = 0.28) in one-play) that, although insignificant ($t[159] = 1.16$, ns) is in the same direction of the effect observed in previous studies (e.g. Keren & Wagenaar, 1987; Keren, 1991).

In summary, then, the current experiment demonstrates the robustness of the known properties of description-based decisions. Thus, it is safe to conclude that the results observed in Experiments 2–4 do not reflect the unique taste of the Technion students; rather, these results reflect the unique properties of feedback-based decisions.

Another interesting finding involves the complex relationship between the three conditions. In only one of the five problems does the maximization rate in the feedback-based conditions fall between the maximization rates in the other two conditions. The feedback-based maximization is highest in two problems and minimal in the other two. Moreover, in two of the conditions the feedback-based results are closer to the one-play results than to the repeated-play results. It is important, however, to recall that the observation that the effect of feedback differs from the effect of repeated play does not prove that the two effects are independent. It is possible that the effect of feedback is an outcome of an interaction between the effect of repeated play and the source of information. Additional research is needed to evaluate this possibility.

BASIC PHENOMENA AND A SIMPLE MODEL

An attempt to summarize the results observed in the eleven feedback-based conditions reveals that the main patterns can be described as a product of two basic psychological tendencies. First, as noted by Thaler et al., the low maximization rate in Problem 1 is naturally explained by *loss aversion*. As in description-based decisions, decision makers behave as if losses loom larger than gains. Interestingly, the aggregate behavior in all the other problems can be explained as a result of *reliance on recent outcomes*.⁹ In Erev and Barron (2003) we quantify these principles within a relatively complex model designed to capture a wide set of learning phenomena. To avoid unnecessary complexity and to highlight the relationship of the current findings to prospect theory, the current section presents a simpler abstraction of these principles. The proposed abstraction is a generalization of Sarin and Vahid's (2001) Payoff Assessment Model.

A value assessment (VA) model

Thaler et al. (1997) accounted for the existence of loss aversion in feedback-based decisions with the addition of two assumptions to prospect theory. The first assumption implies that feedback is used to estimate the payoff distributions of the different alternatives. The second assumption states that decision makers are myopic: they treat each of the decision problems individually (rather than considering the outcome of a series of decisions).

The current results can be easily captured within this framework with a minimal refinement of the first assumption. Specifically, it is sufficient to assume that the weighted value of the different prospects is estimated based on a weighted adjustment process (see March, 1996; Sarin & Vahid, 2001) in which recent outcomes are overweighted. The model can be summarized by the following assumptions:¹⁰

⁹The assertion that recent outcomes are overweighted is consistent with a wide set of studies (see e.g. Duncker, 1935; Messick & Liebrand, 1995; Gigerenzer & Goldstein, 1999). In the current setting reliance on small samples (Kareev, 2000) is likely to lead to a similar effect.

¹⁰Readers interested in deriving the prediction of the model with less effort can use the program available on our web site at <http://techunix.technion.ac.il/~barron/va.htm>. In addition, the program's source code has been made readily accessible for those wishing to construct a simulation themselves.

Adjusted value

The adjusted value of gamble j at trial $t + 1$ is:

$$A_j(t + 1) = (1 - w_t)A_j(t) + (w_t)v(x_t)$$

where $v(x_t)$ is the subjective value of the obtained payoff x_t , and $0 < w_t < 1$ determines the weight of this value. The initial value, $A_j(1)$, is the expected subjective value from random choice; for example, in Problem 4 $A_j(1) = 0.5v(3) + 0.5[0.5v(4) + 0.5v(0)]$. The function $v(\cdot)$ and the weight w_t are defined below.

Recency and exploration

A recency effect can be simply captured with the constraint that adjustment speed (the weight w_t) does not depend on t . However, this abstraction implies high sensitivity to the possibility of sequences of bad outcomes. For example, in Problem 4 a sequence of ‘0’ outcomes could eliminate future H choices, and in Problem 6 a sequence of ‘-4’ outcomes could eliminate future L choice. The similarity of the maximization rates in Problems 4 and 6 suggests that the possibility of bad sequences is not so influential. To address this observation, the current model distinguishes between exploration and exploitation trials.

Each experimental trial is assumed to be an ‘exploration’ or an ‘exploitation’ trial. In exploration trials, the decision maker randomly selects among the alternatives. The probability that trial t is an exploration trial decreases with experience. It is

$$P(t \text{ is an exploration trial}) = \kappa / (t + \kappa)$$

where κ is an exploration strength parameter. The weight w_t (defined in the first assumption) is assumed to depend on whether trial t belongs to the set of ‘exploration trials’. The weight is higher in exploration trials:

$$w_t = \begin{cases} \alpha & \text{if } t \text{ is an exploration trial} \\ \beta & \text{otherwise} \end{cases}.$$

where $0 \leq \beta \leq \alpha < 1$. Thus, giving a higher weight to recent outcomes is captured with the assertion that α and β do not decrease with t .

Loss aversion

As in prospect theory, loss aversion is captured by the subjective value:

$$v(x_i) = \begin{cases} x_i & \text{if } x_i \geq 0 \\ \lambda x_i & \text{if } x_i < 0 \end{cases}$$

Choice rule in exploitation trials

In exploitation trials (trials that do not belong to the set of exploration trials), the gamble with the highest adjusted value is selected (random choice is assumed in exploration trials and when the two gambles have the same adjusted value). Notice that with the parameters $\lambda = 1$, $\kappa = 0$, and $\alpha = \beta$, the model is identical to

Sarin and Vahid's (2001) basic payoff assessment rule. Larger λ implies loss aversion, and with exploration trials ($\kappa > 0$ and $\alpha > \beta$) the model captures recency with limited sensitivity to bad sequences.

To examine if this model can capture the current results, we tried to find parameters that minimize the mean squared deviations between the model's predictions and the observed maximization rates presented in Figures 1–4. However, since the data include only one condition in which loss aversion is relevant, we did not try to estimate λ . Rather, we used the estimate $\lambda = 2.25$ derived by Tversky and Kahneman (1992). The mean squared distance (MSD) between the data and the predictions was minimized when the three free parameters were set to $\kappa = 20$, $\alpha = 0.3$, and $\beta = 0.01$.¹¹

The model's predictions with these parameters are presented in the second column of Figures 1–4. The figures reveal that the model reproduces the main trends: It produces the loss aversion observed in Experiment 1, the reversed certainty in Experiment 2, the reversed payoff domain in Experiment 3, and the overweighting of small probabilities in Experiment 4. Following experience (trial 101 to the end of the experiment) the correlation between the observed and predicted proportions (one value per condition) is 0.94, the MSD score is 0.0015.

To see the intuition behind this good fit, note that the fast adjustment in exploration trials implies a tendency toward random choice when the payoff variability is high, and a tendency to underweight rare events. The combination of fast adjustment and high payoff variability leads toward random choice because the gamble with the higher weighted value is expected to vary from trial to trial. A tendency of this type captures the difference between Problems 2 and 3.

Underweighting of rare events is predicted because these events tend to be underrepresented in the small samples of overweighted recent outcomes. This observation can explain the reversed certainty effect (Experiment 2), the reversed payoff domain effect (Experiment 3) and the results of Experiment 4.

Shortcoming and interpretation

It is important to emphasize that we do not suggest that the model is entirely accurate. Important inaccuracies are related to the observed between-subject variation, and the long-term trends. As suggested by the large standard deviation scores reported in Experiments 1–4, analysis of individual participants shows bimodal (U-shaped) distributions. An extreme example is provided in Problem 11: Of the 24 subjects, 12 selected H in fewer than 20% of the trials, and 5 selected H in more than 80% of the trials. The value assessment model predicts large between-subject variance, but not so large. To produce bimodal distributions, the assumption of identical parameters to all subjects will have to be relaxed.

The failure in capturing the nonlinear long-term learning trends cannot be easily addressed with the value assessment model. To address these trends we (in Erev & Barron, 2002, following Erev et al., 1999) proposed a model that assumes reinforcement learning among cognitive strategies (RELACS). This model captures, for example, the V-shaped curve in Problem 1 because the agents are assumed to slowly learn to stop using a 'loss avoidance' strategy.¹²

Given these shortcomings it is important to recall why the model can be useful. We present the model to provide a clear summary of the main experimental trends. This summary highlights the fact that the unique properties of feedback-based decisions can be a result of a tendency to rely on recent outcomes. Obviously,

¹¹To estimate the parameters, computer simulations were run that derive the models predictions given a wide set of parameter combinations. The estimated parameters are the members of the combination that minimizes the MSD score. Notice that the low value of β reflects the fact that higher values imply high sensitivity to bad sequences and inaccurate prediction of a large difference between Problems 4 and 6.

¹²Interestingly, in addressing binary-choice (two-arm bandit) problems RELACS' advantage is not translated to a much better MSD score. Over the 40 binary-choice conditions analyzed in Erev and Barron, VA and RELACS have almost identical scores (and much better scores than all the other models considered by Erev and Barron). RELACS has a large advantage over VA in addressing problems with large number of choice alternatives.

this tendency has little effect on description-based decisions in which all the information is presented before the decision makers and recency plays no role. Thus, the difference between description- and feedback-based decisions can be explained with the assertion that reliance on recent outcomes is only relevant in the second class of decision problems.

SUMMARY AND POTENTIAL IMPLICATIONS

The current research demonstrates that the study of small feedback-based decision problems can lead to interesting observations. One class of nontrivial observations implies that, in certain settings, immediate feedback does not lead toward expected value maximization. Three types of deviations from maximization are summarized above. A second class of interesting findings involves the direction of the observed deviations. In three cases the deviations in feedback-based decisions are in the opposite direction of the deviations from maximization found in mainstream decision research that focuses on description-based decisions. Experience in feedback-based decisions was found to lead to (1) a reversed common ratio/certainty pattern, (2) a reversed payoff domain effect (more risk aversion in the loss than the gain domain), and (3) an underweighting of small probabilities. Only one of the examined properties of description-based decisions, loss aversion, seems to robustly emerge from feedback.

While these results could reflect a complex and difficult pattern to predict, the current analysis suggests that the opposite is true. In retrospect it is easy to see that the deviations from maximization in feedback-based decisions can be understood as a product of two known psychological principles: loss aversion (Kahneman & Tversky, 1979) and reliance on recent outcomes. Thus, it is easy to capture the results with a simple model.

Implications for static models of choice behavior

It is important to re-emphasize that the current findings do not represent violations of prospect theory and similar static descriptive models. Prospect theory was proposed to capture description-based decisions in important one-shot tasks, so the fact that it does not account for small feedback-based decisions does not refute it. The only implication of the current results for static descriptive models concerns their boundaries. The large differences between description-based and small feedback-based decisions suggest that the generalization of these models to feedback-based decisions requires some modifications. One possible modification is provided above.

Practical implications and future research

At a first glance, the current study could be criticized on the grounds that it focuses on an extremely artificial and unimportant set of decision tasks. According to this criticism, decision makers never face the same important decision problem twice. Obviously, however, this criticism ignores the observation that many common activities involve small feedback-based decisions. And, as noted in the introduction, these small decisions can be consequential even when the expected payoffs of the different alternatives are small and similar. The expected cost of stopping at a 'dark' orange traffic light, for example, is only few seconds.

In order to demonstrate that the study of small feedback-based decisions can lead to important implications, we chose to conclude with a discussion of three examples of empirical phenomena that appear to reflect the unique properties of feedback-based decisions documented here. The first and most important problem involves the usage of safety devices. There is wide agreement in the applied literature (see review

in Grindle et al., 2000) that people under-use safety devices even when this risky behavior cannot be justified as rational. Interestingly, in many cases people voluntarily buy the safety devices that they do not use. For example, in Israel a large proportion of buyers of new car radios choose radios with detachable panels. These radios cost more but appear to be more attractive because they can protect the owners against the small probability event of radio theft. However, most buyers (including the authors of the current paper) detach the panel in the first few days after buying it, but stop detaching it within a short period. Thus, the typical owners behave as if they give much higher weight to the small probability of theft in the 'one-shot-like' buying decision than in the 'feedback-based-like' detaching decision. Presumably, feedback has 'taught' the owners that detaching the panel is a nuisance, while nothing appears to happen when the panel is not detached.¹³

A second example involves the design of casino slot machines (see Haruvy et al., 2001). The common explanations for gambling are based on models of one-shot description-based decisions. According to these explanations, gambling is the result of a limited risk seeking segment in the utility function (Friedman & Savage, 1948), or the overweighting of small probabilities (Kahneman & Tversky, 1979). While these explanations appear to be consistent with certain state lotteries, they cannot account for the popularity of medium-sized prizes in casino slot machines. Under the current model the medium prizes are useful (to the casino owners) because they increase payoff variance and slow the gamblers from learning to stop wasting money.

A third example involves crime prevention. Popular static models of choice behavior imply that people are more sensitive to the magnitude than to the probability of the punishment (see Perry et al., 2002). This prediction is a result of the assumed underweighting of small probabilities (as in prospect theory with the parameters estimated by Tversky & Kahneman, 1992) or risk aversion (in Becker's, 1968, expected utility-based analysis). While this prediction may be correct for major crimes, the current results suggest that the opposite is expected in minor crimes (like running a red light), when the prospective criminals can use feedback. Empirical (e.g. Varma & Doob, 1998) and experimental results (Perry et al., 2002) support the current predictions.

We believe that these and similar applications demonstrate the value of experimental analysis of small feedback-based decisions. This analysis can significantly increase the scope of applications of behavioral decision-making research.

APPENDIX: INSTRUCTIONS (EXPERIMENT 1, PROBLEM 1)

In this experiment you are operating a money machine. Upon pressing a button, you will win or lose a number of points. Your goal is to complete the experiment with as many points as possible.

It is given that there is a difference between the buttons.

The basic payment is 5 Shekels. Your final payment is composed of the points you earn (200 points = 10 Agorot) and the basic payment.

For your information, the exact 'machine' is likely to differ between participants. Good luck.

¹³Some of the respondents bought radios with detachable panels more than once. This observation suggests that adaptation in small decisions is not always transferred to big decisions. The detachable panel problem can also be explained as an example of time inconsistency in consumer behavior (see e.g. Hoch & Loewenstein, 1991). Indeed, it is not clear that the two explanations are inconsistent. The current results can be used to refine models of time inconsistency; they suggest an increase in a particular type of time inconsistency with experience.

The experimental screen:



Note: The third paragraph (conversion rate) differed between experiments and conditions as dictated by the design. The Hebrew words on the screen are ‘Agorot’ (4 agorot = 1 US cent) and ‘total’.

ACKNOWLEDGEMENTS

This research was supported by a grant from the USA–Israel Binational Science Foundation, by the Technion VPR fund, and by the Max Wertheimer Minerva Center. Part of this research was conducted when Erev was a visiting professor at Columbia Business School. We thank Frank Yates, Daniel Kahneman, Ernan Haruvy, Al Roth, and Drew Fudenberg for constructive comments to earlier drafts and presentations, and Dennis Paoli for editorial assistance.

REFERENCES

- Allais, M. (1953). Le comportement de l’homme rationel devant le risque, critique des postulats et axiomes de l’école américaine. *Econometrica*, *21*, 503–546.
- Barkan, R., Zohar, D., & Erev, I. (1998). Accidents and decision making under uncertainty: a comparison of four models. *Organizational Behavior and Human Decision Processes*, *74*, 118–144.
- Barron, G. (2000). The effect of feedback on decision making under uncertainty and risk: the predictive value of descriptive models. Master’s Thesis, Technion-Israel Institute of Technology.
- Barron, G., & Erev, I. (2001). *VA Excel macro for repeated binary choice tasks*. Retrieved July 15, 2002, from <http://techunix.technion.ac.il/~barron/va.htm>
- Becker, G. S. (1968). Crime and punishment: an economic approach. *Journal of Political Economy*, *76*, 169–217.
- Benartzi, S., & Thaler, R. (1995). Myopic loss aversion and the equity premium puzzle. *Quarterly Journal of Economics*, *110*, 75–92.
- Birnbaum, M. H., & Navarrete, J. B. (1998). Testing descriptive utility theories: violations of stochastic dominance and cumulative independence. *Journal of Risk and Uncertainty*, *17*, 49–78.
- Blincoe, L. J. (1994). The Economic Cost of Motor Vehicle Crashes. *National Highway Traffic Safety Administration Technical Report*.
- Böckenholt, U., & Wedell, D. H. (1990). Moderation of preference reversals in the long run. *Journal of Experimental Psychology: Human Perception and Performance*, *16*(2), 429–438.

- Budescu, D. V., Erev, I., Wallsten, T. S., & Yates, J. F. (1997). Introduction to a special issue on stochastic and cognitive models of confidence. *Journal of Behavioral Decision Making*, *10*, 153–155.
- Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, *100*, 432–459.
- Camerer, C. F., & Ho, T. (1999). Experience-weighted attraction in games. *Econometrica*, *64*, 827–874.
- Chen, Y., & Tang, F. (1999). Learning and incentive compatible mechanisms for public goods provision: an experimental study. *Journal of Political Economy*, *106*, 633–662.
- Cheung, Y. W., & Friedman, D. (1998). A comparison of learning and replicator dynamics using experimental data. *Journal of Economic Behavior and Organizations*, *35*, 263–280.
- Cooper, D., Garvin, S., & Kagel, J. (1997). Signalling and adaptive learning in an entry limit pricing game. *Rand Journal of Economics*, *28*, 662–683.
- Daniel, T. E., Seale, D. A., & Rapoport, A. (1998). Strategic play and adaptive learning in sealed bid bargaining mechanism. *Journal of Mathematical Psychology*, *42*, 133–166.
- Duncker, K. (1935). Learning and insight as employed in goal-reaching. *Acta Psychologica, The Hague*, *1*, 77–82.
- Erev, I. (1998). Signal detection by human observers: a cutoff reinforcement learning model of categorization decisions under uncertainty. *Psychological Review*, *105*, 280–298.
- Erev, I., & Barron, G. (2003). On adaptation, maximization, and reinforcement learning among cognitive strategies. Technion working paper.
- Erev, I., & Roth, A. (1998). Prediction how people play games: reinforcement learning in games with unique mixed strategy equilibrium. *American Economic Review*, *88*, 848–881.
- Erev, I., Bereby-Meyer, Y., & Roth, A. E. (1999). The effect of adding a constant to all payoffs: experimental investigation, and implications for reinforcement learning models. *Journal of Economic Behavior and Organizations*, *39*, 111–128.
- Friedman, M., & Savage, L. J. (1948). The utility analysis of choice involving risk. *Journal of Political Economy*, *56*, 279–304.
- Gigerenzer, G., & Goldstein, D. G. (1999). Betting on one good reason: The Take The Best heuristic. In G. Gigerenzer, P. M. Todd, & the ABC Group (Eds.), *Simple heuristics that make us smart*. New York: Oxford University Press.
- Gneezy, U., & Potters, J. (1997). An experiment on risk taking and evaluation periods. *Quarterly Journal of Economics*, *112*, 631–645.
- Gonzalez, R., & Wu, G. (1999). On the shape of the probability weighting function. *Cognitive Psychology*, *38*, 129–166.
- Grindle, A. C., Dickinson, A. M., & Boettcher, W. (2000). Behavioral safety research in manufacturing settings: a review of the literature. *Journal of Organizational Behavior Management*, *20*, 29–68.
- Haruvy, E., Erev, I., & Sonsino, D. (2001). The medium prizes paradox: evidence from a simulated casino. *Journal of Risk and Uncertainty*, *22*, 251–261.
- Herrnstein, R. J., Loewenstein, G. F., Prelec, D., & Vaughan, W., Jr. (1993). Utility maximization and melioration: internalities in individual choice. *Journal of Behavioral Decision Making*, *6*, 149–185.
- Hoch, S. J., & Loewenstein, G. F. (1991). Time-inconsistent preferences and consumer self-control. *Journal of Consumer Research*, *17*, 492–507.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: an analysis of decision under risk. *Econometrica*, *47*, 263–291.
- Kareev, Y. (2000). Seven (indeed, plus or minus two) and the detection of correlations. *Psychological Review*, *107*, 397–402.
- Keren, G. (1991). Calibration and probability judgments: conceptual and methodological issues. *Acta Psychologica*, *77*, 217–273.
- 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.
- Lopes, L. L. (1981). Decision making in the short run. *Journal of Experimental Psychology: Human Learning & Memory*, *7*(5), 377–385.
- Lopes, L. L. (1996). When time is of the essence: averaging, aspiration, and the short run. *Organizational Behavior and Human Decision Processes*, *65*, 179–189.
- MacDonald, D. H., Kagel, J. H., & Battlio, R. C. (1991). Animals' choices over uncertain outcomes: further experimental results. *Economic Journal*, *101*, 1065–1084.
- March, J. G. (1996). Learning to be risk averse. *Psychological Review*, *103*, 309–319.
- Messick, D. M., & Liebrand, W. B. G. (1995). Individual heuristics and the dynamics of cooperation in large groups. *Psychological Review*, *102*, 131–145.
- Myers, J., Suydam, M. M., & Gambino, B. (1965). Contingent gains and losses in risk-taking situations. *Journal of Mathematical Psychology*, *2*, 363–370.

- Perry, O., Haruvy, E., & Erev, I. (2002). Frequent delayed probabilistic punishment in law enforcement. *Economics of Governance*, 3, 71–86.
- Sarin, R., & Vahid, F. (2001). Predicting how people play games: a simple dynamic model of choice. *Games and Economic Behavior*, 34, 104–122.
- Selten, R., & Buchta, J. (1998). Experimental sealed bid first price auctions with directly observed bid function. In D. Budescu, I. Erev, & R. Zwick (Eds.), *Games and human behavior: Essays in honor of Amnon Rapoport* (pp. 53–78). Mahwah, NJ: Erlbaum.
- Thaler, R. H., Tversky, A., Kahneman, D., & Schwartz, A. (1997). The effect of myopia and loss aversion on risk taking: an experimental test. *Quarterly Journal of Economics*, 112, 647–661.
- Thorndike, E. (1898). Animal intelligence: an experimental study of the associative processes in animals. *Psychological Monographs*, 2, 1–109.
- Tversky, A., & Bar-Hillel, M. (1983). Risk: the long and the short. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 9(4), 713–717.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 9, 195–230.
- Varma, K. N., & Doob, A. N. (1998). Deterring economic crimes: the case of tax evasion. *Canadian Journal of Criminology*, 40, 165–184.

Authors' biographies:

Greg Barron is a PhD candidate and lecturer at the Technion Institute of Technology. His research interests are in the areas of learning, decision making, computer simulation and experiments. Other interests include group dynamics, technical analysis, and applied cyber-psychology.

Ido Erev is an Associate Professor of Psychology at the Technion. His current research focuses on modeling and deriving the implications of human adaptation to economic incentives. Examples of implications include the design of rule enforcement and bonus systems, training methods, marketing strategies, and pricing.

Authors' address:

Greg Barron and **Ido Erev**, Faculty of Industrial Engineering and Management, Technion, Haifa, Israel.