

Loss Aversion, Diminishing Sensitivity, and the Effect of Experience on Repeated Decisions[†]

IDO EREV*, EYAL ERT and ELDAD YECHIAM

Max Wertheimer Minerva Center for Cognitive Studies, Faculty of Industrial Engineering and Management, Technion, Haifa, Israel

ABSTRACT

Three experiments are presented that explore the assertion that loss aversion and diminishing sensitivity drive the effect of experience on choice behavior. The experiments are focused on repeated choice tasks where decision makers choose repeatedly between alternatives and get feedback after each choice. Experiments 1a and 1b show that behavioral tendencies that were previously interpreted as indications of loss aversion in decisions from experience are better described as products of diminishing sensitivity to absolute payoffs. Experiment 2 highlights a nominal magnitude effect: A decrease in the magnitude of the nominal payoffs eliminates the evidence for diminishing sensitivity. These and related previous results can be captured with a model that assumes reliance on small samples of subjective experiences, and an increase in diminishing sensitivity with payoff variability. Copyright © 2008 John Wiley & Sons, Ltd.

KEY WORDS myopic loss aversion; prospect theory; learning; reflection effect; explorative sampler

INTRODUCTION

According to the loss aversion hypothesis (Kahneman & Tversky, 1979), the disutility of a loss is larger than the utility of an equivalent gain. Empirical research suggests that this hypothesis captures an important property of the effect of experience on choice behavior. For example, Benartzi and Thaler (1995) rely on the loss aversion hypothesis to explain why many investors have not learned to prefer stocks over bonds even after more than a century in which the average return of stocks was much larger than that of bonds (Mehra & Prescott, 1985).¹ According to this explanation, bonds are preferred because they eliminate the risk of

*Correspondence to: Ido Erev, Max Wertheimer Minerva Center for Cognitive Studies, Faculty of Industrial Engineering and Management, Technion, Haifa 32000, Israel. E-mail: erev@tx.technion.ac.il

[†]Part of this research was conducted when Ido Erev was a Marvin Bower Fellow at the Harvard Business School.

¹Mehra and Prescott call this phenomenon “The Equity Premium Puzzle.” Benartzi and Thaler (1995) explain this puzzle with myopic loss aversion (MLA) which is a combination of two behavioral concepts: myopia (the tendency to evaluate outcomes frequently) and loss aversion. Note that both concepts are equally necessary in explaining the puzzle.

(subjectively) costly losses. Another interesting example is provided by Camerer, Babcock, Loewenstein, and Thaler's (1997) analysis of the behavior of taxi drivers in New York City. This analysis suggests a loss aversion explanation to the observation that drivers tend to work more hours on bad days when the per-hour wage is low but quit earlier on good days in which the wage per-hour is high; a behavioral pattern that contradicts the prediction of the standard theory of labor supply. The authors suggest that the drivers set their reference point on the daily income target and act as if they are loss averse by trying to minimize the possibility of falling short of that reference point.²

However, direct experimental tests of the loss aversion hypothesis lead to contradictory conclusions. Whereas Thaler, Tversky, Kahneman, and Schwartz (1997; and see Barron & Erev, 2003) found deviations from maximization that can be explained by the loss aversion hypothesis, the results reported by Katz (1964) show no evidence of loss aversion.

The main goal of the current study is to improve our understanding of the descriptive value of the loss aversion hypothesis in decisions from experience: contexts where decision makers are not presented with information about the possible outcomes and their likelihoods but have to rely on personal experience.³ In order to achieve this goal we start with an analysis of the problems studied by Thaler et al. (1997).

MIXED GAMBLES AND MIXED RESULTS

Thaler, Tversky, Kahneman, and Schwartz (1997; and see Gneezy & Potters, 1997) examined the role of loss aversion in a simplified stock market. Their basic condition, referred to here as "Mixed," included 200 independent trials. In each trial, the participants were asked to allocate 100 tokens between two assets: A safe bond and a risky stock. Investment in the bond always resulted in a nonnegative outcome. Investment in the stock increased the expected return by a factor of 4, but was associated with high variability and frequent losses. The decisions were made from experience: the participants did not receive any description of the relevant payoff distributions, and had to rely solely on their feedback that was presented graphically⁴ after each trial. The results reveal that the (low expected value) bond attracted about 60% of the investments. To confirm that the attractiveness of the bond reflected loss aversion (rather than risk aversion), Thaler et al. added the "Gain" condition. This condition was identical to the mixed condition, except that a constant was added to all payoffs to eliminate the possibility of losses. In support of the loss aversion hypothesis, this addition increased the attractiveness of the stock.

Barron and Erev (2003) ran a simplified replication of Thaler et al.'s study using a "clicking paradigm": In each of the 200 trials of their study, the participants were asked to select between (click on one of) two unmarked keys (instead of investing tokens). Each selection was rewarded with a draw from the key's payoff distribution. As in the original study the participants did not receive a description of the different distributions, but had to base their decisions on the feedback they received from previous choices. The feedback included a numerical presentation of the obtained payoff. Two problems were compared. Problem "Mixed" was a replication of the mixed condition in Thaler et al., while Problem "Gain" was a variant of the gain condition. The exact payoff distributions in these problems are presented below:

²This suggestion was recently criticized by Ferber (2005).

³Previous research have showed that decisions from experience are different from one-shot decisions that are based on descriptions of the prospects' outcomes and likelihoods (e.g., Barron & Erev, 2003; Hertwig, Barron, Weber, & Erev, 2004).

⁴That is, the returns for "bonds" and "stocks" were presented in bar graphs.

Problem Mixed (Barron & Erev, 2003, following Thaler et al., 1997)

- S A draw from a truncated (at zero) normal distribution with a mean of 25 and standard deviation of 17.7 (implied mean of 25.63) P(S) = 0.70
- R A draw from a normal distribution with a mean of 100 and standard deviation of 354

Problem Gain (Barron & Erev, 2003, following Thaler et al., 1997)

- S A draw from a normal distribution with a mean of 1225 and standard deviation of 17.7 P(S) = 0.49
- R A draw from a normal distribution with a mean of 1300 and standard deviation of 354

The results replicated the pattern observed by Thaler et al. Over the 200 trials the choice rate of the safer, low-expected-payoff, prospect (S) was 70% in Problem Mixed (when R was associated with frequent losses), and only 49% in Problem Gain. In order to clarify the relationship of their results to the loss aversion hypothesis, Thaler et al. used a simplified and myopic variant of prospect theory (Kahneman & Tversky, 1979). The simplification implied by this model involves the assumption of a linear weighting function. The added myopic term asserts that the decision makers consider one decision at a time (rather than considering the payoff distribution implied by a sequence of decisions).

Specifically, Thaler et al. (1997) assumed that choice behavior reflects an attempt to maximize expected subjective value, and the subjective value of outcome x is given by prospect theory's value function. That is,

$$sv(x) = \begin{cases} x^\alpha & \text{if } x \geq 0 \\ -\lambda(-x)^\beta & \text{if } x < 0 \end{cases} \quad (1)$$

Under prospect theory the parameter $\lambda > 1$ captures loss aversion and parameters α and β capture (assuming $0 < \alpha, \beta < 1$) diminishing sensitivity to increases in the absolute payoffs. According to the diminishing sensitivity assumption the subjective impact of a change in the absolute payoff decreases with the distance from zero (see Tversky & Kahneman, 1992, and motivating observations in Stevens, 1957). As noted by Thaler et al. the high S rate in Problem Mixed is predicted by the assertion of a strong loss aversion (high λ). For example, with $\lambda = 2.25$ (and $\alpha = \beta = 1$), the expected subjective values in Problem Mixed are approximately -21 from R, and $+26$ from S. With these parameters, the model implies that R is much more attractive in Problem Gain.

The loss aversion explanation of the pattern discovered by Thaler et al. has many attractive features. It is clear, simple, sufficient, and it clarifies the relationship of the results to a wide set of phenomena that can be naturally explained with the loss aversion hypothesis. However, the loss aversion assertion is not necessary. The same pattern can be captured with diminishing sensitivity. This is the case even under prospect theory (Kahneman & Tversky, 1979), the model used by Thaler et al.: when α is low, S is more attractive in the mixed condition even without loss aversion (i.e., with $\lambda = 1$). Under this "diminishing sensitivity" explanation, S is more attractive in the mixed problem because all the payoffs with the same payoff-sign seem similar. For example, with $\alpha = \beta = 0.5$ (and $\lambda = 1$), the expected subjective values in Problem Mixed are 4.4 from R, and 4.9 from S. With these parameters the model implies similar expected subjective values from R and S in Problem Gain. Notice that Simon's (1955) step-level "satisficing" utility function is an extreme version of the current diminishing sensitivity hypothesis.

Thaler et al.'s selection of the loss aversion explanation was justified by the usage of prospect theory with the parameters estimated by Tversky and Kahneman (1992): $\alpha = \beta = 0.88$, $\lambda = 2.25$. With these parameters the results are driven by loss aversion. However, there are good reasons to doubt the generality of these parameters to the current context. First, many estimations of prospect theory parameters yielded lower α values. For example, Camerer and Ho's (1994) data imply⁵ $\alpha = 0.37$, and Wu and Gonzalez's (1996) data

⁵Camerer and Ho (1994) did not report α , but Wu and Gonzalez repeated their estimation procedure using their data and found $\alpha = .37$ (see Footnote 12 in Wu & Gonzalez, 1996).

imply α values around 0.5. A second and more important reason is the observation that the loss aversion explanation is inconsistent with previous studies of choice behavior in repeated choice tasks. A clear violation of this explanation is provided by Katz (1964). Katz's study included 400 trials. In each trial the participants were asked to guess which of two light bulbs (S or R) would be turned on. The two bulbs were equally likely to be on. Guessing S was safer: The implied payoff was +1 if the guess was correct and -1 otherwise. Guessing R was riskier: The implied payoff was +4 if the guess was correct and -4 otherwise. The participants received no prior information concerning the relevant probabilities, but had to rely on the feedback they received after each trial. The implied choice problem is

Problem Katz (Katz, 1964)

- S +1 with probability 0.5 $P(S) = 0.49$
 -1 otherwise
 R +4 with probability 0.5
 -4 otherwise

The loss aversion assertion that losses loom larger than equivalent gains implies that most people should avoid the larger loss (of -4) and prefer Option S. In violation of this prediction the participants were indifferent between the two options. Notice that Katz's results can be captured by the diminishing sensitivity hypothesis; this hypothesis implies random choice in Katz's Problem. In addition, the results can be captured with a refinement of the loss aversion hypothesis that entails aversion to the possibility of losing (see Erev & Barron, 2005; Erev, Bereby-Meyer, & Roth, 1999).

EXPERIMENT 1A: LOSS AVERSION OR DIMINISHING SENSITIVITY?

The main goal of Experiment 1a was to compare the loss aversion and the diminishing sensitivity explanations of Thaler et al.'s results. We employed the basic clicking paradigm used by Barron and Erev (2003) to replicate Thaler et al.'s results, and focused on the following Problems:

Problem 1 (Mixed)

- S 0 with certainty
 R +1000 with probability 0.5
 -1000 otherwise

Problem 2 (Gain)

- S 1000 with certainty
 R 2000 with probability 0.5
 0 otherwise

Note that in Problem 1 (Mixed) choosing the safer option eliminates the probability of losses. Therefore, the loss aversion hypothesis predicts a higher proportion of S choices in Problem 1 (Mixed) than in Problem 2 (Gain). According to this hypothesis the association of Option R with frequent losses in Problem 1 will decrease its attractiveness. The diminishing sensitivity hypothesis predicts the opposite pattern: random choice in Problem 1 and a strong preference to select S in Problem 2. This is because in Problem 1 the possible gain and loss are of the same distance (1000) from the reference point (0) and thus cancel each other out. In Problem 2, however, the diminishing sensitivity hypothesis implies that the subjective value of the even chance to win 2000 or nothing is reduced at a higher rate than the subjective value of the sure gain of 1000. As control conditions, Experiment 1 also examines Problems 3 and 4 (presented in Table 1): Both hypotheses imply a higher rate of S choices in Problem 3 than in Problem 4.

Experimental design and procedure

The participants in the experiment were 45 Technion students. The experiment used a within-participant design. Each participant was seated in front of a personal computer and was presented with each of the

Table 1. The problems studied in Experiment 1a and the aggregate results

Basic problems	Verbal description	Notation	Proportion of S (safe) choices	
			Mean (<i>SD</i>)	Median
1				
S	0 with certainty	0	0.51	0.50
R	1000 with probability 0.5 −1000 otherwise	(1000, 0.5; −1000)	(0.29)	
2				
S	1000 with certainty	1000	0.70	0.74
R	2000 with probability 0.5 0 otherwise	(2000, 0.5; 0)	(0.21)	
	<i>Mixed-Gain (MG) Score</i>		−0.19 (0.33)	−0.16
3				
S	400 with certainty	400	0.76	0.82
R	1400 with probability 0.5 −600 otherwise	(1400, 0.5; −600)	(0.21)	
4				
S	1400 with certainty	1400	0.66	0.64
R	2400 with probability 0.5 400 otherwise	(2400, 0.5; 400)	(0.23)	
	<i>Mixed-Gain (MG) Score</i>		0.10 (0.27)	0.08

The left-hand column presents the four basic problems studied in Experiment 1a in which participants chose repeatedly between a safer option (S) and riskier option (R). The right-hand columns present the main results over the 100 trials run in the two conditions.

four problems presented in Table 1 for a block of 100 trials. Participants were told that the experiment would include several independent sections, in each of which they would operate a different “computerized money-machine” with two buttons for an unspecified number of trials. Each section involved a repeated play of one of the four problems. In each trial the participants were asked to select one of the buttons. Each selection followed with a presentation of its outcome in points (a draw from the relevant distribution). For example, a selection of Gamble R in Problem 1 (Mixed) resulted in a random draw from a binomial distribution that pays +1000 with probability of 0.5 and −1000 otherwise. This outcome appeared on the selected key and was added to the “accumulated earnings” score. The participants were told that their goal was to maximize their earnings. The points accumulated during the experiment were converted to cash at the rate of 0.01 Agarot (0.0023 US cent) per 1 point. Final payoffs ranged between 26 Sheqels (\$6.19) and 30 Sheqels (\$7.14). The whole procedure lasted about 40 minutes.

The participants received a description of the conversion from points to cash, but did not receive prior information concerning the process that generates the payoff in points (the games’ payoff structure), nor were they informed in advance that the experiment included four sections of 100 trials each (see a translation of the instructions in Appendix A). Before each section they were simply notified that a new game was about to start. They were also told that the new game differed from the previous games. Thus, the participants had to rely on their obtained feedback: the realized payoffs after each choice.

The order of the problems was randomized over participants. The assignment of alternatives to buttons was randomly determined for each participant at the beginning of each section and was fixed during the section.

Results

The right-hand column in Table 1 presents the proportion of S choices over the 100 trials in each of the four problems. A comparison of the proportion of S choices in Problems 1 and 2 reveal that the safer option was *less* popular when it eliminated the probability of losses (51% in Problem 1) than when losses were not possible (70% in Problem 2). In order to evaluate the significance of this pattern we calculated for each participant the difference between the proportion of S choices in the Mixed problem (Problem 1) and the proportion of S choices in the Gain problem (Problem 2). This difference was denoted the Mixed-Gain (hereafter referred to as "MG") score. The mean MG score was -0.19 ($SD = 0.33$), indicating a significant "reversed loss aversion" tendency, $t(44) = -3.85$, $p < .0005$. This result is predicted by the diminishing sensitivity hypothesis, and contradicts the predictions of the loss aversion hypothesis. Additional support for the diminishing sensitivity hypothesis comes from the examination of Problems 3 and 4. In these problems, the safer option tended to be more popular in the mixed problem (76% in Problem 3) than in the gain problem (66% in Problem 4). This difference is also significant (mean MG score = 0.10 , $SD = 0.27$, $t(44) = 2.51$, $p < .02$).

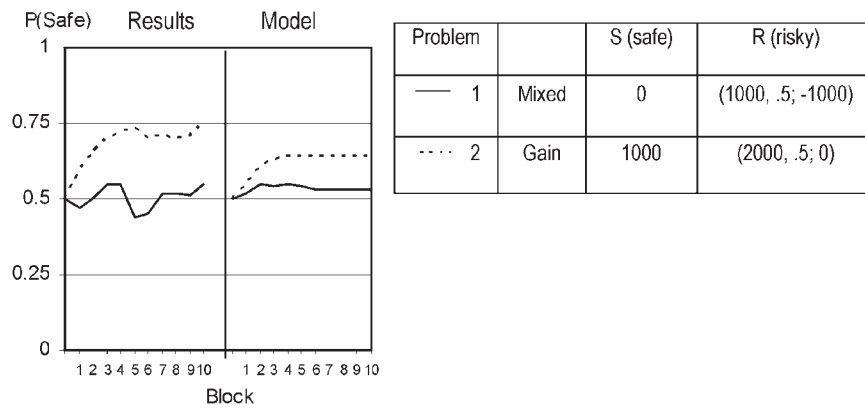
Figure 1 presents the learning curves in these problems (the predictions of the explorative sampler model to be discussed below are presented on the right column of the graph). The results show that the difference between the mixed and gain conditions in each pair of problems increases over time.

An examination of the four problems' order of presentation does not reveal a consistent effect, $F(3, 176) = 1.75$, ns. Nevertheless, to examine the robustness of the difference between Problem 1 and 2 we performed a between-participant analysis that avoids the risk of an order effect. This analysis focuses on the first problem presented to the participants. The observed S rates are 34% in Problem 1 ($n = 11$), and 69% in Problem 2 ($n = 11$). The difference is significant, $t(20) = -3.18$, $p < .005$. Thus, the between-participant analysis agrees with the within-participant analysis.

Another interesting order-related analysis involves the possibility of a "house money" effect and/or a "break even" effect (see Thaler & Johnson, 1990) which assert that decision makers' risk taking is affected by past gains and losses. The "house money" effect asserts that risk taking is facilitated by previous gains. According to the "break even" effect risk taking is facilitated by previous losses but only under the possibility of eliminating these losses. A generalization of these effects to the current setting implies that the accumulated payoffs shift the relevant reference point, and this shift can increase risk taking. The house money effect can be used to predict more risk seeking after gains, and the break even effect can be used to predict more risk seeking after losses. In order to evaluate the house money effect we compared the behavior of the 11 participants who faced Problem 1 first (without house money) with the behavior of the other 34 participants who gained significant amounts before facing Problem 1. The results show no indication of a house money effect: Participants took more risk without house money (66% over the 11 participants than faced Problem 1 first) than with house money (44% over the remaining 34 participants; $t(43) = 2.30$, $p < .03$, for the difference between the two groups).

In order to evaluate the break even effect we re-analyzed the behavior of the participants who started the experiment with Problem 1. The analysis focused on the decisions that were made in trials 11–100. Seven of the 11 participants experienced both negative accumulated payoffs and accumulated payoffs of zero during these 90 trials. Only one of these seven participants exhibited more risk seeking given negative accumulated payoffs. The other six took more risk while the accumulated payoff was 0. Thus, they did not exhibit the break even effect.

a Problems 1 & 2:



b Problems 3 & 4:

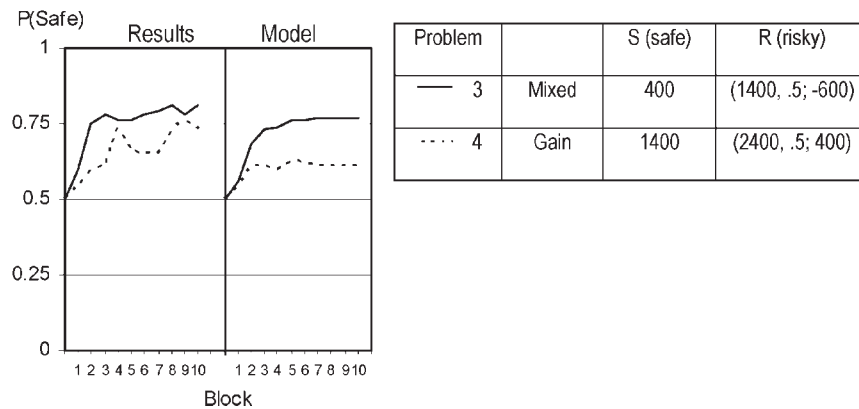


Figure 1. Proportion of safe choices in 10 blocks of 10 trials and the sensitive sampler model’s predictions in each of the four problems studied in Experiment 1a. The notation $(X, p; Y)$ refers to a gamble that yields a payoff of x with probability p and y otherwise.

To summarize, the results support the diminishing sensitivity hypothesis and contradict the loss aversion hypothesis. This pattern does not appear to be a product of a house money and/or break even effects. The clearest evidence against the loss aversion assumption is provided in Problem 1 in which half the participants preferred the risky option despite the high loss involved with this option.

EXPERIMENT 1B: DIMINISHING SENSITIVITY OR ZERO AVOIDANCE?

According to one interpretation of the results of Experiment 1a, the relative low choice rate of the safe alternative in Problem 1 reflects an attempt to avoid the payoff “0”. This interpretation is consistent with previous demonstrations that boredom can facilitate risk taking in laboratory tasks (e.g., Lei, Noussair, & Plott, 2001). Specifically, one could speculate that participants were bored by repeatedly getting nothing and

this facilitated their risk taking. Experiment 1b evaluates this zero avoidance hypothesis by focusing on the following pair of problems:

Problem 5 (Mixed)

- S +200 with probability 0.5 $P(S) = 0.43$
 -200 otherwise
 R +1000 with probability 0.5
 -1000 otherwise

Problem 6 (Gain)

- S 1200 with probability 0.5 $P(S) = 0.72$
 800 otherwise
 R 2000 with probability 0.5
 0 otherwise

The current problems differ from Problems 1 and 2 of Experiment 1a in that the safer alternative (S) is a gamble. Thus, if the pattern observed in Experiment 1a (more S choices in Problem 2 compared to Problem 1) is a product of zero avoidance then the difference between the problems should decrease. The zero avoidance hypothesis predicts higher S rate in Problem 5 than in Problem 1. The diminishing sensitivity hypothesis, however, predicts that the differences in risk taking between the current problems will be similar to the differences observed between Problems 1 and 2. Specifically, it predicts 50% risk taking in Problem 5 (because the magnitude of the positive and negative values is symmetrically distributed around zero), and a strong preference to S in Problem 6 since the value of 2000 in R is discounted to a higher degree than the values of the safer alternative.

Experimental design and procedure

The participants in the experiment were 30 Technion students. The current experiment used the same design and procedure as the first experiment with the exception that it was focused on problems 5 and 6, presented above.

Results

The right-hand column in Table 2 presents the proportion of S choices over the 100 trials in each of the two problems. A comparison of the proportion of S choices in Problems 5 and 6 reveals that the safer option was *less* popular when it reduced losses (43% in Problem 5) than when losses were not possible (72% in Problem 6). The mean MG score was -0.28 ($SD = 0.40$), which implies a significant “reversed loss aversion” tendency, $t(29) = -3.92$, $p < .0005$. This result is predicted by the diminishing sensitivity hypothesis and cannot be explained with the zero avoidance hypothesis. Additional support for the diminishing sensitivity hypothesis comes from the observation that the proportion of risk taking in Problem 5 was not significantly different than 50%, $t(29) = -1.34$, NS.

Figure 2 presents the learning curves in the Problems 5 and 6 (the predictions of the explorative sampler model are presented in the right column of the graph). Similarly to the learning curves in Experiment 1a, the difference between the two conditions increases over time.

As in Experiment 1a, no order effects were found in the current experiment, $F(1, 28) = 0.48$, NS. Analyzing only the first ordered problems does not change the pattern of results: the proportion of safe choices remains lower in Problem 5 (45%) than in Problem 6 (64%).

EXPERIMENT 2: THE NOMINAL MAGNITUDE EFFECT

The diminishing sensitivity hypothesis, supported above, appears to be inconsistent with the observation that the main regularities documented in previous studies of decisions from experience can be captured with

Table 2. The problems studied in Experiment 1b and the aggregate results

Basic problems	Verbal description	Notation	Proportion of S (safe) choices	
			Mean (SD)	Median
5				
S	200 with probability 0.5 -200 otherwise	(200, 0.5; -200)	0.43	0.47
R	1000 with probability 0.5 -1000 otherwise	(1000, 0.5; -1000)	(0.28)	
6				
S	1200 with probability 0.5 800 otherwise	(1200, 0.5; 800)	0.72	0.81
R	2000 with probability 0.5 0 otherwise	(2000, 0.5; 0)	(0.31)	
	<i>Mixed-Gain (MG) Score</i>		-0.28 (0.40)	-0.34

The left-hand column presents the two basic problems studied in Experiment 1b in which participants chose repeatedly between a safer option (S) and riskier option (R). The right-hand columns present the main results over the 100 trials run in the two conditions.

Problems 5 & 6:




Figure 2. Proportion of safe choices in 10 blocks of 10 trials and the explorative sampler model's predictions in each of the pair of problems studied in Experiment 1b. The notation $(X, p; Y)$ refers to a gamble that yields a payoff of x with probability p and y otherwise.

models that assume risk neutrality (see Erev & Barron, 2005). One possible resolution of this inconsistency is based on the observation that most of the problems considered by Erev and Barron involve low nominal payoffs, while the problems studied above involve high nominal payoffs. Under this explanation, the apparent inconsistency reflects a nominal magnitude effect: On average, decision makers exhibit risk neutrality when the nominal values are low, and they behave as if their value function is S shaped (i.e., risk aversion in the gain domain, and risk seeking in the loss domain) when the nominal payoffs are high.⁶

⁶A similar argument was suggested by Holt and Laury (2002). They noted that relative risk aversion tends to increase with higher stakes. According to the current hypothesis diminishing sensitivity (that implies risk aversion in the gain domain) increases with nominal payoffs (even when the actual stakes do not change).

The current experiment examined this “nominal magnitude” hypothesis by studying the four problems presented in Table 3 under two “nominal magnitude” conditions. The left-hand column in Table 3 presents the basic version of the four problems. Under Condition Low, the feedback after each choice was a draw from the distribution presented in the basic gamble column in Table 3. Condition High was identical to Condition Low except that the payoffs in points were multiplied by a hundred, and the conversion rate from points to money was divided by 100. In other words, the nominal values in the high condition differed in two orders of

Table 3. The problems studied in Experiment 2 and the aggregate results

Basic problems	Verbal description	Notation	Proportion of S (safe) choices				
			Condition Low		Condition High		
			Mean (SD)	Median	Mean (SD)	Median	
7							
S	A draw from the interval (0,1)	$0 + u(0,1)$	0.49 (0.33)	0.45	0.57 (0.28)	0.62	
R	A draw from the interval (-1,0) with probability 0.5 A draw from the interval (2,3) otherwise	$(-1, 0.5; 2) + u(0,1)$					
8							
S	A draw from the interval (3,4)	$3 + u(0,1)$	0.55 (0.25)	0.52	0.47 (0.24)	0.47	
R	A draw from the interval (2,3) with probability 0.5 A draw from the interval (5,6) otherwise	$(2, 0.5; 5) + u(0,1)$					
9							
S	A draw from a truncated (at zero) normal distribution with a mean of 0.25 and standard deviation of 0.177 (implied mean of 0.256)	TN(0.25, 0.177, 0)	0.49 (0.30)	0.47	0.60 (0.26)	0.64	
R	A draw from a normal distribution with a mean of 1 and standard deviation of 3.54	N(1, 3.54)					
10							
S	A draw from a truncated (at 12) normal distribution with a mean of 12.25 and standard deviation of 0.177 (implied mean of 12.256)	TN(12.25, 0.177, 12)	0.53 (0.26)	0.53	0.50 (0.26)	0.51	
R	A draw from a normal distribution with a mean of 13 and standard deviation of 3.54 <i>Mixed-Gain Score (MG)</i>	N(13, 3.54)					
			-0.05 (0.27)	-0.01	0.10 (0.22)	0.11	

The left-hand column presents the four basic problems studied in Experiment 2 in which participants chose repeatedly between a safer option (S) and riskier option (R). The right-hand columns present the main results over the 100 trials run in the two conditions: In Condition Low the presentation of payoffs on the experimental screen matched the verbal description in the left side of the table. In Condition High these payoffs were multiplied by 100.

magnitude from the nominal values in the Low condition. Thus, the typical payoff in Condition Low involved several points, and the typical payoff in Condition High involved several hundred points. Nevertheless, the two conditions were economically identical: The sole objective difference between the two point magnitude conditions was the addition of a decimal point to the feedback in Condition Low.

In order to evaluate the robustness of the results, Problems 7 and 8 were associated with bimodal distributions and Problems 9 and 10 were associated with normal distributions. Note that in Problems 7 and 9 the risky option is associated with frequent losses, whereas the safer option is not. Following Thaler et al. (1997), Problems 8 and 10 (the “gain” problems) were created by the addition of a constant to Problems 7 and 9 (the “mixed” problems), respectively.

The nominal magnitude hypothesis predicts a difference between the two conditions. Specifically, it predicts that the proportion of S choices in the mixed problems (7 and 9) will be higher than in the gain problems (8 and 10) in Condition High but not in Condition Low.

Experimental design and procedure

Experiment 2 compared two between-participant groups (i.e., High and Low point magnitudes). Each group faced the four problems presented in Table 3, using a within-participant design: Each participant was faced with each of the four problems for a block of 100 trials. One hundred Technion students who did not participate in the first experiment, served as paid participants in the current study. Fifty were randomly assigned to Condition Low, and the other 50 were assigned to Condition High. The procedure was identical to Experiment 1, with the exception that the current study focuses on the problems presented in Table 3. The instructions were as in Experiment 1. The participants did not know the payoff distributions but they were told the exact rate in which their points would be converted to money. The conversion rates in this experiment were 2 agorot (about 0.46 US cent) per 1 point in Condition Low and 0.02 agorot (about 0.0046 US cent) per 1 point in Condition High. Final payoffs ranged between 32 Sheqels (\$7.61) and 39 Sheqels (\$9.28).

Results

The right-hand columns in Table 3 present the (mean and median) proportion of S choices over the 100 trials in each of the four problems under the two conditions. The results reveal a clear nominal magnitude effect. In Condition Low the safer option was slightly *less* popular in the mixed problems when it eliminated the probability of loss (49% in Problem 7, and 49% in Problem 9) than in the gain problems (55% in Problem 8, and 53% in Problem 10). The mean MG score was -0.05 ($SD = 0.27$). This difference is not significantly different than 0, and it reflects no evidence for the pattern implied by the diminishing sensitivity effect.


In Condition High, however, the safer option tended to be more popular in the mixed problems when it eliminated the probability of loss (57% in Problem 7, and 60% in Problem 9) than in the gain problems (47% in Problem 8, and 50% in Problem 10). In this condition the mean MG score was 0.10 ($SD = 0.22$). This difference is significant, $t(49) = 3.34$, $p < .002$.

Comparison of the two conditions reflects the pattern predicted by the nominal magnitude hypothesis. The mean MG score in Condition High (0.10) is significantly higher than the mean MG score in Condition Low (-0.05 ; $t(98) = 3.05$, $p < .003$).

The learning curves are presented in Figure 3 (together with the predictions of the explorative sampler model). They show that the pattern described above is robust to experience. Indeed, the difference between the two point magnitude conditions slightly increases over time.

3.a Problems 7 & 8:

Problem		S (safe)	R (risky)
— 7	Mixed	$0+u(0,1)$	$(2, .5;-1)+u(0,1)$
···· 8	Gain	$3+u(0,1)$	$(5, .5;2)+u(0,1)$



3.b Problems 9 & 10:

Problem		S (safe)	R (risky)
— 9	Mixed	$TN(0.25, 0.177, 0)$	$N(1.00, 3.54)$
···· 10	Gain	$TN(12.25, 3.54, 12)$	$N(13.00, 3.54)$




Figure 3. Proportion of safe choices under high and low point magnitudes, and the explorative sampler model's predictions in 10 blocks of 10 trials in each of the four problems studied in Experiment 2. The left- and right-hand columns present the results and the predictions of the sensitive sampler model respectively in each of the point magnitude conditions (High and Low). The notation $(X, p; Y)$ refers to a gamble that yields a payoff of x with probability p and y otherwise. The notation $u(V, Z)$ refers to a draw from a uniform distribution between v and z . The notation $TN(B, F, P)$ refers to a draw from a normal distribution with mean of B , standard deviation of F , and the payoff is truncated at point P .

A QUANTITATIVE SUMMARY AND ALTERNATIVE ABSTRACTIONS

Experiment 2 takes one step toward relating the current results to previous experimental studies of decisions from experience; it shows that the apparent difference between Experiment 1 and the previous experimental studies of decisions from experience, reviewed by Erev and Barron (2005), can be a product of a nominal magnitude effect. The main goal of the current section is to take another step in the same direction; it tries to refine the models proposed by Erev and Barron in order to capture the current findings. Specifically, it searches for a model that can capture the diminishing sensitivity and payoff magnitude effects as well as the main regularities considered by Erev and Barron with a single set of parameters.

The models proposed by Erev and Barron were designed to address two robust deviations from maximization (of expected payoffs) that were not considered above. The first deviation is the payoff variability effect (see Myers & Sadler, 1960): High payoff variability reduces sensitivity to payoff difference. The second deviation can be described as underweighting of rare (low probability) events (see Barron & Erev, 2003; Yechiam & Busemeyer, 2005). Erev and Barron's analysis shows that both deviations can be captured with the assertion that decisions from experience are driven by best reply to small samples of experiences (see related assumptions in Biele, Erev, & Ert, 2007; Erev, Glozman, & Hertwig, 2008; Hertwig et al., 2004; Hochman & Erev, 2007; Kareev, 2000; Osborne & Rubinstein, 1998). According to this assertion the decision maker recalls a small set of past experiences with each alternative, and selects the alternative that is associated with the better average experiences in the recalled set. Lebiere, Gonzalez, and Martin (2007) extend this analysis and show the value of a model that assumes similarity-based weighting. This model implies high sensitivity to small sample of experiences (experiences in "similar" situations) and lower sensitivity to other experiences.

We believe that the diminishing sensitivity pattern, documented above, can be used to improve our understanding of the effect of small samples on choice behavior. Specifically, we hypothesize that the addition of the diminishing sensitivity assumption to models that assume oversensitivity to small samples, can improve the value of these models. The current section evaluates this optimistic hypothesis by considering variants of the explorative sampler model described below. The analysis focuses on the effect of the abstraction of diminishing sensitivity on the model's fit of the 14 experimental conditions considered here and the results summarized by Erev and Barron (2005).

The explorative sampler model

The model can be summarized with the following assumptions:

A1: Exploration and exploitation. The agents are assumed to consider two cognitive strategies: exploration and exploitation (see Denrell & March, 2001; Gittins, 1979 for discussions of the value of this distinction).

Exploration implies a random choice. The probability of exploration is 1 in the very first trial, and depends on the availability of information concerning the forgone payoffs in the following trials. When this information is available the probability of exploration in trial $t > 1$ is $0 < \varepsilon < 1$ —a free parameter. When information concerning the forgone payoffs is not available, the probability of exploration reduces toward an asymptote (at ε) with experience. The effect of experience on the probability of exploration depends on the number of trials (T) in the experiment. Exploration diminishes quickly when T is small, and slowly when T is large.⁷ This assumption is quantified as follows:

$$P(\text{Explore}_t) = \varepsilon^{(t-1)/(t+T^\delta)} \quad (2)$$

where δ is a free parameter that captures the sensitivity to the length of the experiment.

⁷Implicit in this abstraction is the simplification assumption that the decision makers know the value of T . This assumption is incorrect, but it is likely to provide good approximation of the estimated number. Since the decision makers know the expected length of the experiment in minutes, and the number of subsections, it is natural to assume that they can develop a good estimate of T .

A2: Experiences and sampling. The experiences with each alternative include the set of observed outcomes yielded by this alternative in previous trials. In addition, when feedback is limited to the obtained payoff, the subjective value of the very first outcome is recalled as an experience with all the alternatives.

Under exploitation the agent draws (with replacement) a sample of m_t past experiences with each alternative. All previous experiences are equally likely to be sampled. The value of m_t at trial t is assumed to be randomly selected from the set $\{1, 2, \dots, \kappa\}$ where κ is a free parameter.

The sampling algorithm is assumed to depend on the available information. When feedback is limited to the obtained payoffs the sampling from the experiences with the different alternatives is independent. When the foregone payoffs are known (the decision makers receive complete feedback that includes the payoff from the unselected alternatives), the distinct samples are perfectly correlated. The decision maker selects one set of m_t trials, and the outcomes in those trials are used to determine the values of the different alternatives.

A3: Regressiveness, diminishing sensitivity, and choice. The recalled subjective values of the outcome x from selecting alternative j at trial t is assumed to be affected by two factors: regression to the mean of all the experiences with the relevant alternative (in the first $t - 1$ trials), and diminishing sensitivity. Regression is captured with the assumption that the regressed value is $R_x = (1 - w)x + (w)A_j(t)$, where $0 < w < 1$ is a free parameter and $A_j(t)$ is the average outcome from the relevant alternative.⁸

Diminishing sensitivity is captured with a variant of prospect theory's (Kahneman & Tversky, 1979) value function that assumes

$$sv(x) = \begin{cases} R_x^{\alpha_t} & \text{if } R_x \geq 0 \\ -(-R_x)^{\alpha_t} & \text{if } R_x < 0 \end{cases} \quad (3)$$

where $\alpha_t = (1 + V_t)^{-\rho}$, $\rho \geq 0$ is a free parameter, and V_t is a measure of payoff variability. V_t is computed as the average absolute difference between consecutive obtained payoffs in the first $t - 1$ trials (with an initial value at 0). The parameter ρ captures the effect of diminishing sensitivity: large ρ implies quick increase in diminishing sensitivity with payoff variability.

The estimated subjective value of each alternative at trial t is the mean of the subjective value of the alternative's sample in that trial. Under exploitation the agent selects the alternative with the highest estimated value.

Estimation and results

In order to evaluate the model, we simulated virtual replications of the 14 conditions described above and the 40 conditions reviewed by Erev and Barron (2005, see Table 4⁹). The simulated participants arrived at their choices on the basis of the model's assumptions. Thus, we can compare the choice proportions predicted by the model to the empirically observed choice proportions. The following five steps were taken in each trial:

1. The trial decision mode (exploration or exploitation) was determined (using Equation 2). If the selected type was exploration, one option was randomly chosen and the process moved to step 3.

⁸Implicit in this regressiveness (the assumption $0 < w < 1$) is the assumption, introduced by Lebiere et al. (2007), that all the experiences are weighted (because all the experiences affect the mean). This implicit assumption is necessary to capture a thought experiment in which the decision maker chooses between "1000 with certainty" and the gamble "1001; 9; 0 otherwise."

⁹The 40 conditions were run under three experimental paradigms. Under the Probability Learning (PL) paradigm the decision maker is asked to predict which of two mutually distinctive events will occur in the next trial, and can see when the trial ends which event occurred. Under the Minimal Information (MI) paradigm the individual is asked during every trial to select one of two unmarked buttons, and gets feedback concerning the payoff of the chosen button. The Complete Feedback (CF) paradigm is similar to the Minimal Information paradigm with the exception that the decision maker is presented with the values of both buttons after each choice, but her payoff is determined by the selected button.

Table 4. Comparison of the results and the predictions of the explorative sampler model

Experiment	Problem	R (risky)	S (safe)	Proportion of S (safe) choices	
				Observed	Explorative sampler
1a					
0	1	(1000, 0.5; -1000)	0	0.51	0.47
	2	(2000, 0.5; 0)	1000	0.72	0.73
+400	3	(1400, 0.5; -600)	400	0.75	0.73
	4	(2400, 0.5; 400)	1400	0.65	0.60
1b					
±200	5	(1000, 0.5; -1000)	(200, 0.5; -200)	0.43	0.51
	6	(2000, 0.5; 0)	(1200, 0.5; 800)	0.72	0.73
2					
Low	7L	(2, 0.5; -1) +u(0,1)	0 + u(0,1)	0.49	0.44
	8L	(5, 0.5; 2) +u(0,1)	3 + u(0,1)	0.55	0.40
	9L	N(1.00, 3.54)	TN(0.25, 0.177) truncated at 0	0.49	0.45
	10L	N(13.00, 3.54)	TN(12.25, 0.177) truncated at 12	0.53	0.43
High	7H	(200, 0.5; -100) +u(0,100)	0 + u(0,100)	0.62	0.61
	8H	(500, 0.5; 200) +u(0,100)	200 + u(0,100)	0.52	0.43
	9H	N(100, 354)	TN(25, 17.7) truncated at 0	0.60	0.55
	10H	N(1300, 354)	TN(1225, 354) truncated at 1200	0.50	0.44

Notice that the letter added to the problem number in Experiment 2 reflects the nominal magnitude condition: L for Low, and H for High. The Explorative Sampler column refers to the predicted proportion of S choices under the Explorative Sampler model.

2. The following actions were taken in the case of exploitation

- a. The sample size used by the agent (m_t) was determined.
- b. A sample of m_t outcomes was drawn with replacement from the experience with each alternative.
- c. The alternative with the higher mean subjective value in the sample was selected.

3. The outcomes were realized by drawing from the objective payoff distributions.

4. The experiences (observed outcomes) were stored.

5. The measures of the payoff variance (V_t), and the average payoffs ($A_j(t)$) were updated.

The model's parameters were estimated (using a grid search method with mean squared deviation criteria) to simultaneously fit the 54 experimental conditions. The estimated parameters were $\rho=0.15$, $w=0.3$, $\delta=0.55$, $\varepsilon=0.08$, and $\kappa=6$. Tables 4 and 5 and the right panels in Figures 1–3 present the implied predictions. These exhibits show that the model reproduces the main patterns of the results. Specifically, the model reproduces the three very different effects of the addition of a constant to all the payoffs. As in the human data this addition: (i) increases the rate of choosing Safe in the upper panel of Figures 1 and 2, (ii) decreases the rate of choosing Safe in the lower panel on Figure 1 and under the high condition in Figure 3, and (iii) has little effect under the Low condition in Figure 3. In addition to the qualitative

Table 5. Comparison of the results and the explorative sampler model predictions in the problems studied by Erev and Barron (2005)

Problem and paradigm	Number of trials	Alternative H	Alternative L	Proportion of (L) choices	
				Observed	Explorative sampler
1 MI	200	(11)	(10)	0.16	0.08
2 MI	200	(11)	(1, 0.5; 19)	0.32	0.38
3 MI	200	(1, .5; 21)	(10)	0.42	0.56
4 MI	200	(-10)	(-11)	0.07	0.08
5 MI	200	(-10)	(-1, .5; -21)	0.46	0.48
6 MI	200	(-1, .5; -19)	(-11)	0.52	0.43
7 CF	200	(11)	(10)	0.06	0.04
8 CF	200	(1, .5; 21)	10	0.37	0.45
9 CF	200	-10	-11	0.06	0.04
10 CF	200	-10	-1, .5; -21	0.45	0.45
Problems 11-14: G is the gamble (-5, .5; +5)					
11 MI	200	$N(21,3)$	$N(18,3)$	0.15	0.20
12 MI	200	$N(21,3)$	$N(18,3) + G$	0.24	0.28
13 MI	200	$N(21,3) + G$	$N(18,3)$	0.41	0.30
14 MI	200	$N(21,3) + G$	$N(18,3) + G$	0.38	0.32
Problems 15-20: H = (x if E; -x if not-E), L = (-x if E; x if not-E)					
15 PL	400	$x = 1, p(E) = .6$		0.40	0.31
16 PL	400	$x = 10, p(E) = .6$		0.38	0.30
17 PL	400	$x = 1, p(E) = .7$		0.23	0.20
18 PL	400	$x = 10, p(E) = .7$		0.20	0.20
19 PL	400	$x = 1, p(E) = .8$		0.23	0.21
20 PL	400	$x = 10, p(E) = .8$		0.12	0.12
21 MI	400	(4, .8; 0)	(3)	0.36	0.40
22 MI	200	(4, .2; 0)	(3, .25; 0)	0.48	0.46
23 MI	200	(32, .1; 0)	(3)	0.71	0.74
24 MI	200	(32, .025; 0)	(3, .25; 0)	0.69	0.68
25 MI	200	(-3)	(-32, .1; 0)	0.60	0.70
26 MI	200	$N(100,354)$	$TN(25, 17.7)$	0.69	0.54
27 MI	200	$N(1300, 354)$	$N(1225, 17.7)$	0.48	0.41
28 MI	200	$N(1300, 17.7)$	$N(1225, 17.7)$	0.16	0.08
Problems 29-35: H = (G if E; B if not-E), L = (B if E; G if not-E)					
29 PL	300	$P(E) = .75, G = 5, B = 0$		0.17	0.15
30 PL	300	$P(E) = .75, G = 5, B = -5$		0.05	0.16
31 PL	500	$P(E) = .7, G = 6, B = 2$		0.20	0.20
32 PL	500	$P(E) = .7, G = 4, B = 0$		0.22	0.20
33 PL	500	$P(E) = .7, G = 2, B = -2$		0.11	0.20
34 PL	500	$P(E) = .7, G = 0, B = -4$		0.13	0.20
35 PL	500	$P(E) = .7, G = -2, B = -6$		0.18	0.20
Problems 36-39: $P(E) = .7; P(F) = .9; H = (G \text{ if } E\&F, B \text{ if not-}E\&F, 0 \text{ otherwise});$ $L = (B \text{ if } E\&F, G \text{ if not-}E\&F, 0 \text{ otherwise})$					
36 PL	550	$G = 6, B = 2$		0.17	0.20
37 PL	550	$G = -2, B = -6$		0.18	0.20

(Continues)

Table 5. (Continued)

Problem and paradigm	Number of trials	Alternative H	Alternative L	Proportion of (L) choices	
				Observed	Explorative sampler
38 MI	550	G = 6, B = 2		0.22	0.27
39 MI	550	G = -2, B = -6		0.22	0.25
40 MI	400	(-3)	(-4, .8; 0)	0.40	0.35

The left-hand columns present the 40 problems studied by Erev and Barron (2005) and the observed results (prop. of L choices in the first 100 trials). The paradigms are: MI = minimal information, CF = complete feedback, and PL = probability learning. The notation $(x, p; y)$ describes a gamble that pays x with probability p , y otherwise. The notation $(x \text{ if } E; y \text{ if not-}E)$ implies a gamble that pays x if E occurs and y otherwise. $N(x, y)$ means a draw from a normal distribution with mean x and standard deviation y , $TN(25, 17.7)$ is a truncated (at zero) normal distribution. Alternative H is associated with higher expected value than Alternative L. The right-hand column presents the prediction of the explorative sampler model.

reproduction, the model provides good quantitative fit of the aggregated results: The mean square deviation (MSD) between the observed and predicted proportions presented in Tables 4 and 5 is 0.0044. This score is similar to the MSD scores of the best models in Erev and Barron’s (2005) analysis. Thus, the refined model is as accurate as these models in capturing Erev and Barron’s data, and outperforms these models in capturing the current data.

Alternative abstractions

Two alternative abstractions of the diminishing sensitivity assertion were examined. The first uses Equation (1)’s abstraction of the diminishing sensitivity assumption. Notice that this abstraction allows loss aversion and assumes constant diminishing sensitivity. The MSD score of the model with this subjective value function is minimized with the parameters $\alpha = 0.7$, $\lambda = 1.3$, $\delta = 0.55$, $\epsilon = 0.08$, and $\kappa = 6$: The MSD score is 0.0087. The relatively high MSD score reflects the fact that Equation (1)’s power value function cannot capture the nominal magnitude effect (and, with the estimated parameters, the model cannot describe the pattern observed in experiment 1).

A second alternative abstraction adds the possibility of loss aversion to Equation (3)’s power function. It assumes Equation (4)’s subjective value function:

$$sv(x) = \begin{cases} R_x^{\alpha} & \text{if } R_x \geq 0 \\ -\lambda(-R_x^{\alpha}) & \text{if } R_x < 0 \end{cases} \tag{4}$$

The analysis of this model shows that the addition of the loss aversion parameter does not improve the fit. Optimal fit is obtained with $\lambda = 1$. Thus the MSD score and the other parameters are identical to those of the simpler model that use Equation (3)’s abstraction.

Shortcomings and possible extensions

The explorative sampler model has three major shortcomings. The first involves the fact that the model under-predicts some of the sequential dependencies in the data. Specifically it under-predicts the tendency to select the alternative with the best recent outcome, and the tendency to repeat recent choices. Biele et al. (2007) show that this limitation can be corrected with the assumption that recency and inertia are products of decision modes that are not modeled here.

A second shortcoming involves the (implicit) assumption of a static environment. The explorative sampling model ignores the possibility that the environment can change. Hochman and Erev (2007; and Biele

et al., 2007) show that this shortcoming can be corrected with the assumption of contingent sampling. Under this assumption the sampling is contingent on an assessment of the state of the environment.

A third shortcoming involves the assumption that the probability of selecting the different modes is independent of the experience. Erev and Barron (2005) highlight one solution to this problem. Under their solution the probability of selecting each mode is determined by a reinforcement learning process (and see related ideas in Rieskamp & Otto, 2006; Stahl, 2000).

Loss aversion and individual differences

The analysis presented above focuses on the behavior of the typical participant. Thus, it suggests that on average, participants are equally explorative of gains and losses, but does not imply that equal sensitivity to gains and losses is general. Indeed, sensitivity to gains compared to losses is at the heart of many of the current theories of individual differences (e.g., Gray, 1994; Higgins, 1997) and of learning models that seek to study decisions at the individual level (e.g., Busemeyer & Stout, 2002; Wallsten, Pleskac, & Lejuez, 2005; Yechiam et al., 2006). The current findings do not contradict these models. What the current findings imply is that across individuals, in the conditions studied here the loss aversion tendency is balanced, so that there are only small differences in the average loss aversion across different individuals.

GENERAL DISCUSSION

The original goal of the current research was to improve our understanding of the effect of loss aversion on decisions from experience. We hoped to propose a refined abstraction of loss aversion that could explain why the effect of experience appears to be sensitive to loss aversion in some settings (Thaler et al., 1997) but not in others (Katz, 1964). The experimental results led us in a different direction: They suggest that the effect of experience in repeated decisions does not appear to reflect loss aversion for the average participant.

The clearest evidence against the hypothesis that the typical decision maker exhibits loss aversion is provided by Problem 1. The typical participant was indifferent between the status quo (payoff of 0) to an equal chance to win 1000 and lose 1000. This result contradicts Kahneman and Tversky's (1979) original definition of loss aversion (losses loom larger than gains), and Erev and Barron's (2005) revised abstraction (an effort to minimize the probability of losses).

In addition, the current results demonstrate that previous findings that were interpreted as evidence for loss aversion in decisions from experience are better described with the assertion of a strong diminishing sensitivity effect. For example, the tendency to prefer safe outcomes that ensure a positive return over risky outcomes with a much higher average return (Barron & Erev, 2003; Thaler et al., 1997) is explained with the assertion of low sensitivity to the difference between the different gains.

Finally, the results suggest that the extent to which decision-makers exhibit the diminishing sensitivity effect is a function of the nominal payoff magnitude. Strong diminishing sensitivity was observed when the feedback involved a gain or loss of hundreds of points, but not when the payoff involved several points. Indeed, when the nominal payoffs were low, the modal behavior exhibited risk neutrality.

The effect of loss aversion on experience in natural settings

The current results appear to be inconsistent with the observation that the loss aversion hypothesis provides an elegant explanation for the effect of experience in many natural decision environments (see Benarzi & Thaler, 1995; Camerer et al., 1997). Under one explanation of this inconsistency, it reflects a difference between "small decisions" (the situations examined here), and bigger decisions (situations in which the absolute difference between the expected values of the different alternatives is high).

A second feasible explanation is based on the observation that there are many alternative explanations for the empirical phenomena interpreted as indications of loss aversion. For example, the suggestion that many individuals are under-invested in the stock market, analyzed by Benartzi and Thaler (1995), can be explained through the diminishing sensitivity hypothesis supported here.¹⁰

We believe that additional research is needed in order to compare the two explanations. The current data cannot be used to determine if the behavioral tendencies observed here are likely to emerge in big decisions.

Loss aversion in decisions under risk, and in riskless choice

Another important disclaimer involves the role of loss aversion in decisions that are based on a description of the relevant payoff distributions. The current results do not question the validity of the loss aversion hypothesis in the context of decisions under risk (one-shot choices among numerically described lotteries), the focus of Kahneman and Tversky (1979).¹¹ Nor do they do so in the context of riskless choice (see Shapira, 1981; Tversky & Kahneman, 1991). The current results do suggest that the loss aversion phenomenon is less general than we originally believed.

One interesting relationship between the current findings and loss aversion studies in other contexts involves the possibility that loss aversion can be a forecasting error (Kermer, Driver-Linn, Wilson, & Gilbert, 2006). According to this assertion people overestimate the impact of potential losses on their actual behavior. The current findings support this assertion by showing that people do not overweight experienced losses compared to gains in their decisions.

Interpretations and implications of the nominal magnitude effect

The nominal magnitude effect, documented in Experiment 2, is consistent with Holt and Laury's (2002) observation that relative risk aversion tends to increase with the magnitude of (nominal) payoffs. These findings seem to challenge recent abstractions of risk taking that assume constant relative risk aversion (e.g., Weber, Shafir, & Blais, 2004).

We believe that the most important implication of the nominal magnitude effect is the identification of a boundary condition for the diminishing sensitivity effect. It implies that when the nominal (and objective) payoffs are low, choice behavior can be captured with a simple model that assumes risk neutrality. Only when the nominal magnitudes are large is the diminishing sensitivity assumption important. The explorative sampler model presents one abstraction of this observation.

Derivation of additional implications requires better understanding of the factors that contribute to this effect. Two factors are likely to be particularly important. The first involves perceptual sensitivity. The perception of large nominal values is more demanding, and likely to involve a stronger perceptual bias.

Obviously, however, the nominal values dimension is only one of many dimensions that can affect perceptual bias. For example, graphical presentation of the outcomes (of the type used by Thaler et al., 1997) is likely to enhance diminishing sensitivity. Another interesting manipulation involves a presentation of foregone payoff (the payoff from the alternative that was not selected). Ert and Erev (in preparation) ran a replication of Experiment 1 with information concerning foregone payoffs. Their results replicate the

¹⁰Notice that there are many natural investment problems in which the loss aversion and diminishing sensitivity hypotheses lead to different predictions. One example involves the choice between an individual stock and an index fund. It is commonly assumed that an index fund is associated with same expected return as an individual stock but with less variability (risk). Thus, as in Katz (1964), loss aversion implies a bias toward the index fund, while the diminishing sensitivity hypothesis implies random choice. Recent research suggests that individual investors deviate from the textbook model in the direction of selecting individual stocks (e.g., Barber & Odean, 2000; Blume & Friend, 1975; Kelly, 1995).

¹¹However, recent research (Ert & Erev, 2007; Schmidt & Traub, 2002) does question the robustness of the loss aversion hypothesis in decisions under risk.

qualitative pattern described above, but reveal smaller differences between the two conditions that can be captured with the assertion that the presentation of foregone payoff decreases diminishing sensitivity.

A second factor involves generalization. It is possible that the presentation of large nominal values increases the tendency to generalize from previous experiences with decisions that involve more substantial amounts of money. This assertion implies a relationship between the nominal magnitude effect and Holt and Laury's (2002) incentive effect.

Potential practical implications: the example of safety rules

An attempt to derive the potential practical implications of the current results reveals two difficulties. First, as suggested above, the current data cannot be used to determine if the behavioral tendencies observed here are likely to emerge in big decisions. Second, the explorative sampler model that captures the main results highlights two behavioral tendencies, "diminishing sensitivity" and "reliance on small samples," that can lead to contradictory predictions.

We believe that these difficulties reduce the set of environments that can be reliably analyzed based on the current results, but they do not eliminate this set; there are many environments in which small decisions (involving small absolute differences between the expected values of the different alternatives) have consequential economic implications. Moreover, in many of these cases, the two tendencies captured by explorative sampler model reinforce each other. For an example consider the value of enforcing safety rules. Specifically, think about situations in which decision makers have to choose between a safe and a riskier action. A concrete example involves a pedestrian (human or chicken) and a road that he, she, or it plans to cross when the pedestrian light is red. This decision maker has to choose between waiting for the green light (the safer option), and crossing during the red light (the riskier option).

The risky option is likely to lead to a gain of few seconds, but there is a small probability that it will lead to a much larger loss. Thus, a naïve generalization of the loss aversion hypothesis suggests that the decision maker is likely to deviate from expected utility maximization in the direction of being "too cautious." The current results lead to the opposite prediction. Diminishing sensitivity implies bias toward risk seeking because it implies insufficient sensitivity to the large potential losses. A bias in the same direction is predicted by the assumed reliance on small samples: Since low probability events are likely to be under-represented in most small samples than the decision maker is likely to behave as if he or she believes that "it won't happen to me." Therefore, the two psychological assumptions abstracted in the explorative sampler model imply a bias towards risk taking in this set of situations. Consequently, the current analysis suggests that the value of the enforcement of safety rules is likely to be larger than the value estimated under the assumption of loss aversion or even rational choice.

Summary

The current analysis rejects the hypothesis that loss aversion drives the effect of experience on repeated decisions. Rather, it suggests that the main behavioral regularities observed for the average participant in previous studies of decisions from experience reflect two robust tendencies: diminishing sensitivity relative to a reference point, and reliance on small samples.

APPENDIX A: THE INSTRUCTIONS FOR THE EXPERIMENTAL TASK

"Hello,

In this experiment you will play a number of different games. In each game you will operate a money machine. Each button press will lead to winning or losing a number of points (depending on the button you

choose). Your goal in the experiment is to win as many points as possible. There could be differences in the number of points produced by each of the buttons. Your final bonus will be determined by the total number of points earned in the game (100 points = 1 Agora). For your information, it is highly likely that the machine would be different for each participant.

Good luck.”

ACKNOWLEDGEMENTS

We thank Greg Barron, Frank Yates, and three anonymous reviewers for their valuable suggestions, and Dana Dahan for running the experiments. This research was supported by a grant from the Israel Science Foundation.

REFERENCES

- Barber, B. M., & Odean, T. (2000). Trading is hazardous to your wealth: The common stock investment performance of individual investors. *Journal of Finance*, *55*, 773–806.
- 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.
- Benartzi, S., & Thaler, R. H. (1995). Myopic loss aversion and the equity premium puzzle. *Quarterly Journal of Economics*, *110*, 73–92.
- Biele, G., Erev, I., & Ert, E. (2007). Learning, risk attitude, and hot stoves in restless bandit problems. MPI, Working Paper.
- Blume, E., & Friend, I. (1975). The asset structure of individual portfolios and some implications for utility functions. *Journal of Finance*, *30*, 585–603.
- Busemeyer, J. R., & Stout, J. C. (2002). A contribution of cognitive decision models to clinical assessment: Decomposing performance on the Bechara gambling task. *Psychological Assessment*, *14*, 253–262.
- Camerer, C. F., & Ho, T. H. (1994). Violations of the betweenness axiom and nonlinearity in probability. *Journal of Risk and Uncertainty*, *8*, 167–196.
- Camerer, C. F., Babcock, L., Loewenstein, G., & Thaler, R. H. (1997). Labor supply of New York City cabdrivers: One day at a time. *Quarterly Journal of Economics*, *112*, 407–441.
- 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.
- 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.
- Erev, I., Glozman, I., & Hertwig, R. (2008). What impacts the impact of rare events, *Journal of Risk and Uncertainty*, *36*, 153–177.
- Ert, E., & Erev, I. (2007). Loss aversion in decisions under risk and the value of a symmetric simplification of prospect theory. Technion, Working Paper.
- Ert, E., & Erev, I. (in preparation). Learning, diminishing sensitivity, and hot stoves in decisions from experience. Technion, Working Paper.
- Ferber, H. S. (2005). Is tomorrow another day? The labor supply of New York City cabdrivers. *Journal of Political Economy*, *113*, 46–82.
- Gittins, J. C. (1979). Bandit processes and dynamic allocation indexes. *Journal of the Royal Statistical Society Series B-Methodological*, *41*, 148–177.
- Gneezy, U., & Potters, J. (1997). An experiment on risk taking and evaluation periods. *Quarterly Journal of Economics*, *112*, 631–645.
- Gray, J. A. (1994). Three fundamental emotion systems. In P. Ekman, & R. J. Davidson (Eds.), *The nature of emotion: Fundamental questions* (pp. 243–247). New York: Oxford University Press.

- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*, 534–539.
- Higgins, T. E. (1997). Beyond pleasure and pain. *American Psychologist*, *52*, 1280–1300.
- Hochman, G., & Erev, I. (2007). The partial reinforcement extinction effect and the role of contingent sampling. Technion, Working Paper.
- Holt, C. A., & Laury, S. K. (2002). Risk aversion and incentive effects. *American Economic Review*, *92*, 1644–1655.
- 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.
- Katz, L. (1964). Effects of differential monetary gain and loss on sequential two-choice behavior. *Journal of Experimental Psychology*, *68*, 245–249.
- Kelly, M. (1995). All their eggs in one basket: Portfolio diversification of U.S. households. *Journal of Economic Behavior and Organization*, *27*, 87–96.
- Kermer, D. A., Driver-Linn, E., Wilson, T. D., & Gilbert, D. T. (2006). Loss aversion is an affective forecasting error. *Psychological Science*, *17*, 649–653.
- Lebiere, C., Gonzalez, C., & Martin, M. (2007). Instance-based decision making model of repeated binary choice. In Proceedings of the 8th International Conference on Cognitive Modeling. Ann Arbor, Michigan, USA.
- Lei, V., Noussair, C. N., & Plott, C. R. (2001). Nonspeculative bubbles in experimental asset markets: Lack of common knowledge of rationality vs. actual irrationality. *Econometrica*, *69*, 831–859.
- Mehra, R., & Prescott E. C. (1985). The Equity Premium: A Puzzle. *Journal of Monetary Economics*, *15*, 145–161.
- Myers, J. L., & Sadler, E. (1960). Effects of range of payoffs as a variable in risk taking. *Journal of Experimental Psychology*, *60*, 306–309.
- Osborne, M. J., & Rubinstein, A. (1998). Games with procedurally rational players. *American Economic Review*, *88*, 834–847.
- Rieskamp, J., & Otto, P. E. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, *135*, 207–236.
- Schmidt, U., & Traub, S. (2002). An experimental test of loss aversion. *Journal of Risk and Uncertainty*, *25*, 233–249.
- Shapira, Z. (1981). Making tradeoffs between job attributes. *Organizational Behavior and Human Performance*, *28*, 331–355.
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, *69*, 99–118.
- Stahl, D. O. (2000). Rule learning in symmetric normal-form games: Theory and evidence. *Games and Economic Behavior*, *32*, 105–138.
- Stevens, S. S. (1957). On the psychophysical law. *Psychological Review*, *64*, 153–181.
- Thaler, R. H., & Johnson, E. J. (1990). Gambling with the house money and trying to break even: The effect of prior outcomes on risky choice. *Management Science*, *36*, 643–660.
- 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.
- Tversky, A., & Kahneman, D. (1991). Loss aversion in riskless choice: A reference dependent model. *Quarterly Journal of Economics*, *107*, 1039–1061.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*, 297–323.
- Wallsten, T. W., Pleskac, T., & Lejuez, C. W. (2005). Modeling a sequential risk taking task. *Psychological Review*, *112*, 862–880.
- Weber, E. U., Shafir, S., & Blais, A. R. (2004). Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychological Review*, *111*, 430–445.
- Wu, G., & Gonzalez, R. (1996). Curvature of the probability weighting function. *Management Science*, *12*, 1676–1690.
- Yechiam, E., & Busemeyer, J. R. (2005). Comparison of basic assumptions embedded in learning models for experience based decision-making. *Psychonomic Bulletin and Review*, *12*, 387–402.
- Yechiam, E., Goodnight, J., Bates, J. E., Busemeyer, J. R., Dodge, K. A., Pettit, G. S., et al. (2006). A formal cognitive model of the go/no-go discrimination task: Evaluation and implications. *Psychological Assessment*, *18*, 239–249.

Authors' biographies:

Ido Erev is the ATIS' Women's Division Professor of Industrial Engineering and Management at the Technion, and a Marvin Bower fellow at HBS (in 2007–2008). His current research focuses on decisions from experience and the economics of small decisions. Examples of implications include the design of rule enforcement systems and training methods.

Eyal Ert is a PhD candidate of Behavioral Sciences at the Technion. His current research interests focus on models of learning and decision making, and their implications to everyday life, consumer behavior, and social interactions.

Eldad Yechiam is a Senior Lecturer at the Industrial Engineering and Management Faculty at the Technion. His research focuses on individual differences in adaptive learning and decision making, and the development of mathematical models that explore these differences.

Authors' addresses:

Ido Erev, Eyal Ert and **Eldad Yechiam**, MaxWertheimer Minerva Center for Cognitive Studies, Faculty of Industrial Engineering and Management, Technion, Haifa, Israel.