LOSS AVERSION, DIMINISHING SENSITIVITY, AND THE EFFECT OF EXPERIENCE ON REPEATED DECISIONS

IDO EREV

EYAL ERT

ELDAD YECHIAM

December 20, 2006

Abstract: Previous studies lead to contradictory conclusions regarding the role of loss aversion in repeated choice tasks. The current paper presents two experiments that explore these inconsistencies. The first study shows that behavioral tendencies that were previously interpreted as indications of loss aversion are better described as products of diminishing sensitivity to absolute payoffs. The second study suggests that the implied diminishing sensitivity may not reflect true risk attitude. Rather, it appears to reflect perceptual sensitivity. These and related results can be captured with a model that approximates a joint generalization of prospect theory (Kahneman and Tversky, 1979) and probability matching.

Address all correspondence to Ido Erev, the Max Wertheimer Minerva Center for Cognitive Studies, Faculty of Industrial Engineering and Management, Technion, Haifa, Israel, erev@tx.technion.ac.il. This research was supported by a grant from the Israel Science Foundation.

I. Introduction

According to the loss aversion hypothesis (Kahneman and 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 human decision-making. For example, Benartzi and Thaler [1995] show that the loss aversion hypothesis can explain why many investors have not learned to prefer stocks over bonds even after 70 years in which the average return of stocks was four times larger than that of bonds. According to this explanation, bonds are preferred because they eliminate the risk of (subjectively) costly losses. Another interesting example is provided by Camerer et al.'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.

However, direct experimental tests of the loss aversion hypothesis lead to contradictory conclusions. Whereas Thaler et al. ([1997]; and see Barron and Erev [2003]) found deviations from maximization that can be explained by the loss aversion hypothesis, the results reported by Katz [1964] show no evidence for 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. In order to achieve this goal we start with an analysis of the problems studied by Thaler et al. [1997].

II MIXED GAMBLES AND MIXED RESULTS

Thaler et al. ([1997]; and see Gneezy and 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 four, but was associated with high variability and frequent losses. The decisions were made from experience: the participants did not receive a description of the relevant payoff distributions, and had to rely on their feedback that was presented graphically[1] 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. In each of the 200 trials of their study, the participants were asked to select between 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:

Problem Mixed (Barron and 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 and 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 variant of prospect theory (Kahneman and Tversky [1979]). Specifically, they 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,

$$(1) \qquad v(x) = \begin{cases} x^{\alpha} & if \quad x \geq 0 \\ -\lambda |x|^{\alpha} & if \quad x < 0 \end{cases}$$

Under prospect theory the parameter α captures a diminishing sensitivity of the decision maker to large absolute payoffs, and the parameter λ captures loss aversion. 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 $\alpha = 1$), 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. Nevertheless, this explanation has shortcomings. The most important shortcoming involves the observation that the loss aversion assertion is not necessary. This is the case even under the model proposed by Thaler et al. When diminishing sensitivity parameter (α) is low, S is more attractive in the Mixed condition even without loss aversion (i.e., with $\lambda = 1$). This is the case because extreme diminishing sensitivity (low α) implies weak sensitivity to the difference between different gains. Under this "diminishing sensitivity" explanation, S is more attractive in the mixed problem because all the positive payoffs look similar (and different from the negative payoffs). For example, with $\alpha = .5$ (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 subjective expected values from R and S in Problem Gain.

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]: $\alpha = .88$, $\lambda = 2.25$. 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 imply[2] $\alpha = .37$, and Wu and Gonzalez's [1996] data imply α values around 0.5. Second, Tversky and Kahneman's estimation of λ was based on a pricing task rather than on a choice task. Finally, and most importantly, the loss aversion explanation is inconsistent with previous studies of choice behavior in repeated choice tasks. A clear violation of this explanation is provided in 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

Loss aversion predicts a preference for 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 implies aversion to the possibility of losing (see Erev and Barron, [2005]).

III. EXPERIMENT 1: LOSS AVERSION OR DIMINISHING SENSITIVITY?

The main goal of Experiment 1 was to compare the loss aversion and the diminishing sensitivity explanations of Thaler et al.'s results. It uses the basic experimental paradigm used by Barron and Erev [2003] to replicate Thaler et al.'s results, and focuses 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 point magnitude (1000) from the reference point (0) and thus cancel each other out. In Problem 2, however, the diminishing sensitivity implies that the subjective utility of the even chance to win 2000 or nothing is reduced at a higher rate than the subjective utility of the sure gain of 1000. As control conditions, Experiment 1 also examines Problems 3 and 4 (presented in Table I): Both hypotheses imply a higher rate of S choices in Problem 3 than in Problem 4.

<Insert Table I>

<u>Experimental Design</u>

The participants in the experiment were 20 Technion students. The experiment used a within-participant design. Each participant was faced with each of the four problems presented in Table I 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. In each trial the participants were asked to select one of the buttons. Each selection followed with a presentation of its outcome (a draw from the relevant distribution). 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, and that at the end of the experiment their accumulated points would be converted to cash at the rate of .01 Agarot (about .00023 US cent) per 1 point. Final payoffs ranged between 26 Sheqels (about $5.90) and 30 Sheqels (about $6.80).

The participants received no prior information about the game's payoff structure. Before each section they were simply told that they were about to start a new game. In Sections 2, 3 and 4, 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 four sections corresponded to the four problems, with the order of the problems being randomized over participants. In each section, the payoffs associated with each button (alternative) were determined by random draws from the corresponding payoff 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. 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 I 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 reveals that the safer option was *less* popular when it eliminated the probability of loss (48% in Problem 1) than when losses were not possible (70% in

Problem 2). In order to evaluate the significance of this pattern we computed an "MG" score for each participant as the difference between proportion of S choices in the Mixed problem (Problem 1) and the proportion of S choices in the Gain problem (Problem 2). The mean MG score was -0.22 (SD = .35), which implies a significant "reversed loss aversion" tendency ($t(19) = -2.82$, $p < .01$). 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 for examination of Problems 3 and 4. In these problems, the safer option tended to be more popular when it eliminated the probability of loss (75% in Problem 3) than when losses were not possible (66% in Problem 4). This difference is significant (mean MG score = 0.09, SD = .22, $t(19) = 1.94$, $p < .035$, one tail).

Figure I presents the learning curves in these problems. The results show that the difference between the two conditions increases over time. For conciseness, the interactions with time are not reported.

< Insert Figure I >

IV. EXPERIMENT 2: RISK ATTITUDE OR PERCEPTUAL SENSITIVITY?

Experiment 2 was designed to compare two distinct interpretations of the diminishing sensitivity hypothesis, supported above. Under one interpretation, the results reflect the decision makers' risk attitude, specifically the attitude toward monetary gains and losses.

Under a second interpretation, the results reflect perceptual sensitivity. This "Perceptual sensitivity" hypothesis allows for the possibility that the decision makers tend to be risk neutral when they are presented with small (and clear) payoffs that are easy to differentiate (e.g., Erev and Barron [2005]). However, this tendency is masked by perceptual bias that can be captured with the assumption that the (perceived) subjective value of a particular payoff is given by an S-shaped function. Thus, perceptual manipulations like an increasing of the point payoff magnitude (like in Barron and Erev [2003]), or a graphical representation of the feedback (like in Thaler et al. [1997]), facilitate an S shape value function since they increase the difficulty to discriminate between the different payoffs.[3]

Experiment 2 compares these hypotheses by studying the four problems presented in Table II under two "point magnitude" conditions. The left-hand column in Table II 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 II. The nominal magnitude in Condition High was similar to that in Barron and Erev [2003]. It 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 a hundred. Notice that the sole difference between the two point magnitude conditions was the addition of a decimal point to the feedback in Condition Low.

<Insert Table II>

In order to evaluate the robustness of the results, Problems 5 and 6 are associated with bimodal distributions and Problems 7 and 8 are associated with normal distributions. Note that in Problems 5 and 7 the risky option is associated with frequent losses, whereas the safer option is not. Following Thaler et al. [1997], Problems 6 and 8 (the "gain" problems) were created by the addition of a constant to Problems 5 and 7 (the "mixed" problems) respectively.

The risk attitude hypothesis asserts that the willingness to take risks is not sensitive to the different presentations of payoffs. This assertion implies similar S rates in the Low and High point magnitude conditions. Under the natural assumption of larger perceptual error in Condition High than in Condition Low, the perceptual sensitivity hypothesis predicts a difference between the two conditions. Specifically, it predicts that the proportion of S choices in the mixed problems (5 and 7) will be higher than in the gain problems (6 and 8) in Condition High but not in Condition Low.

Experimental Design

Experiment 2 compared two between-participant groups (i.e., High and Low point magnitudes). Each group faced the four problems presented in Table 2, 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 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 II. The conversion rates in this experiment were 2 agorot (about .046 US cent) per 1 point in Condition Low and 0.02 agorot (about .00046 US cent) per 1 point in Condition High. Final payoffs ranged between 32 Sheqels (about $7.30) and 39 Sheqels ($8.90).

Results

The right-hand columns in Table II present the (mean and median) proportion of S choices over the 100 trials in each of the four problems studied in Condition Low and Condition High respectively. The results reveal a large point 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 5, and 49% in Problem 7) than in the gain problems (55% in Problem 6, and 53% in Problem 8). The mean MG score was -0.09 (SD = .53). This difference is not significantly different than 0, and it reflects no evidence for diminishing sensitivity.

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 5, and 60% in Problem 7) than in the gain problems (47% in Problem 6, and 50% in Problem 8). In this condition the mean MG score was 0.21 (SD = .44). This difference is significant ($t(49) = 3.34$, $p < .002$).

Comparison of the two conditions reflects the pattern predicted by the perceptual sensitivity hypothesis. The mean MG score in Condition High (0.21) is significantly higher than the mean MG score in Condition Low (-0.09; $t(98) = 3.05$, $p < .003$).

The learning curves are presented in Figure II. They show that the pattern described above is robust to experience. Indeed, the difference between the two point magnitude conditions slightly increases over time.

<Insert Figure II>

V. RELATIONSHIP TO PREVIOUS STUDIES OF DECISIONS FROM EXPERIENCE

Previous experimental studies of decisions from experience (see review in Erev and Barron [2005]) highlight two robust deviations from maximization (of expected payoffs) that are independent of the loss aversion hypothesis considered here. The first deviation is the payoff variability effect (see Myers and 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 and Erev [2003]). Erev and Barron [2005] show that both deviations can be captured with simple models that generalize the probability matching idea (see Estes [1950]) and assume best reply to small samples of experiences. In order to clarify the relationship of the current results to these studies we chose to summarize the results with a model that approximates a joint generalization of Prospect Theory (Kahneman and Tversky [1979]) and probability matching.

Sampling Prospect Theory

The current model, referred to as Sampling Prospect Theory (SPT, see Erev, Glozman and Hertwig [2006]), assumes that the subjective value of the payoff x in environment *e* is given by a generalization of prospect theory's value function[4]:

$$(2) \qquad v(x \mid e) = \begin{cases} x^{\alpha_e} & if \quad x \geq 0 \\ -\lambda |x|^{\alpha_e} & if \quad x < 0 \end{cases}$$

The generalization implied by equation (2) involves the possibility that the environment affects the shape of the function. The term $\alpha_e < 1$ abstracts the diminishing sensitivity in environment *e*, and the parameter λ abstracts the attitude toward losses.

The current results imply two assertions concerning the shape of the value function. First, they suggest that in the environments we considered the value of λ is close to 1. Second, they suggest that the value of α_e depends on the magnitude of the nominal payoff in the relevant environment. The value of α_e appears to increase toward 1 when the nominal magnitude decreases. We quantify this assertion with the following equation:

(3) $$\alpha_e = \alpha + (1-\alpha)^{\beta(Se+1)}$$

The value *Se* in Equation (3) is the standard deviation (in points) of the risky option in environment *e*. Thus, it captures the magnitude of the nominal payoffs. The parameter $\beta \geq 0$ captures the importance of the nominal magnitude effect. The parameter $0 < \alpha < 1$ captures the minimal value of the diminishing sensitivity term. Thus, when $\beta > 1/(Se+1)$, equation (3) implies contingent diminishing sensitivity. Diminishing sensitivity is strong (as α_e approaches its minimum value) when the nominal payoff is large, and weak (as α_e approaches 1) when the nominal payoff is low.

In order to use Equations (2) and (3) to predict behavior, it is necessary to add assumptions concerning the relationship between the subjective value and the final choice. Following Erev and Barron [2005] we assume two decision modes: Exploration and Exploitation. Random choice is assumed during exploration. The probability that a given trial is an exploration trial is assumed to equal ρ -- an exploration parameter. During exploitation, decision maker *i* is assumed to recall a sample of k_i experiences (k_i previous trials) with the feasible alternatives, and to select the alternative with the highest average subjective value in this sample. The exact value of k_i is assumed to be drawn from the set of discrete numbers between 1 and κ, where κ is a sampling size parameter (see related ideas in Kareev [2000]; Osborne and Rubinstein [1998]).

Notice that with the parameters $\lambda = 1$, $\alpha = 1$, $\rho = 0$, and $\kappa = 1$ SPT implies probability matching: The predicted probability that a particular option will be selected equals the probability that this option will lead to the best payoff in a randomly selected trial. To see the relationship of SPT to prospect theory, consider the case: $\rho = 0$, and very large β and κ. In this case the model approximates the predictions of prospect theory under the assumption of a linear weighting function. The approximation results from the fact that α_e converges toward α when β increases, and a very large value for κ implies that the weight of each possible outcome will be identical to its objective probability. By adding the assumption that the probability of each outcome being sampled depends on its rank, SPT can approximate the predictions of the complete version of prospect theory (this possibility is not studied here).

In order to find out whether the SPT model can capture the main results described above, we simulated virtual decision makers for the different experimental conditions. 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 simulation proceeded as follows:

1. The sample size used by agent k_i was drawn from the relevant distribution.

2. The sample of k_i outcomes was drawn from the objective payoff distribution of each of the gambles.
3. The subjective value of each outcome was computed (using Equations 2 and 3).
4. The average subjective values in the samples were recorded.
5. The “best” choice was determined on the basis of the average subjective values in the samples.

We estimated the model's parameters in three ways. The first set of estimates was derived (using a grid search method with mean squared deviation criteria) to fit the 12 experimental conditions studied here under the constraint $\lambda = 1$. The estimated parameters are $\alpha = 0.1$, $\beta = 1.8$, $\rho = 0.6$, and $\kappa = 9$. The mean square deviation (MSD) between the observed and predicted proportions is .0040. The implied predictions are presented in Table III.

<Insert Table III>

Notice that the constraint $\lambda = 1$ implies no loss aversion for the typical participant. Under this constraint the disutility of a loss is equal to the utility of an equivalent gain. To test whether this constraint impairs the fit of the data, we re-estimated the model without it. The grid search considered values of λ from 0.8 to 1.5 with jumps of 0.1, and the estimated parameters did not change. That is, the estimation suggests that $\lambda = 1$ fits the data better than the other values we considered. This result suggests that under the current model the behavior of the typical subject exhibits equal sensitivity to gains and losses.

The third set of parameters was estimated to simultaneously fit the current results and 40 conditions used by Erev and Barron [2005]. Table IV presents the additional conditions. They were selected by Erev and Barron to demonstrate the main behavioral regularities observed in experimental studies of decisions from experience.[5] This analysis is potentially interesting as the SPT model is a variant of the best static 4-parameter model proposed by Erev and Barron to fit their data. The main difference between SPT and that model, referred to as extended probability matching, involves the abstraction of the effect of losses. Thus, under the optimistic assertion that the SPT model is an improvement over Erev and Barron's extended probability matching model, it should capture both data sets with the same sets of parameters. The estimated parameters (with the constraint $\lambda = 1$) are $\alpha = 0.2$, $\beta = 0.1$, $\rho = 0.4$, and $\kappa = 5$. The MSD score is 0.0053 over the 52 conditions (0.0062 over Table III's 12 conditions, and 0.0051 over Table IV's 40 conditions). The extended probability matching model is as accurate as SPT on Erev and Barron's 40 conditions (MSD = 0.0050), but it fails to capture the current data. Its best MSD score over the 52 conditions is 0.0086. Thus, it seems that SPT provides a better summary of the main behavioral regularities.

<Insert Table IV>

Dynamic extensions

One important shortcoming of SPT is implied by its static nature. SPT predicts that the tendency to rely on small samples (which leads to underweighting of rare events) is robust to experience; this prediction does not change even when underweighting of rare events leads to counterproductive outcomes. To clarify this shortcoming, consider decision makers who face the following choice problem for many trials:

Problem 9 (a thought experiment)

S 0 with certainty

R +1 with probability 0.95

-1000 otherwise

With the parameters estimated above (over the 52 conditions) the current model predicts a high rate of R choices (71%). This prediction seems unreasonable. It is natural to assume that human decision makers can learn to prefer the safer and much more attractive option (S) in this example.

We believe that Problem 9's thought experiment implies that the SPT model is too simple. The assumption of a static decision rule has to be relaxed. Our favorite relaxation of this assumption involves the assertion that the SPT rule is only one of several "cognitive strategies" that people consider in a repeated choice environment (see related assertions in Erev and Roth [1998]; Erev and Barron [2005]). Another natural strategy is the fictitious play rule: a selection of the strategy that has led to a higher average payoff in the past (see Brown [1951]). We believe that the probability of selecting each of these strategies can be captured with an assumption of reinforcement learning among the different strategies.

One abstraction of this idea is provided with Erev and Barron's [2005] Reinforcement Learning Among Cognitive Strategies (ReLACS) model. A variant of this model (which uses the current results to refine the abstraction of the effect of losses) outperforms SPT in fitting the current results, and can address the thought experiment in Problem 9.[6] Yet, the advantage of this model over SPT in fitting the current results is small relative to the cost (the added complexity). Thus, we leave the task of comparing this model to other dynamic extensions of SPT to follow-up research.

Loss aversion and individual differences

The analysis presented above focuses on the behavior of the typical participant. Thus, it suggests that typical participants are equally sensitive to 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 difference (e.g., Gray [1994]; Higgins [1997]) and of reinforcement learning models that seek to study decisions at the individual level (e.g., Busemeyer and Stout [2002]; Yechiam et al. [2006]; Wallsten et al. [2005]). 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.

VI. 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. Our participants were indifferent between the status quo (payoff of 0) and 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 strong diminishing sensitivity. For example, the tendency to prefer safe bonds that ensure a positive return over risky stocks with a much higher average return (Thaler et al. [1997]; Barron and Erev [2003]) is explained with the assertion of low sensitivity to the difference between the different gains. This observation implies that the main results can be captured with a simplified version of prospect theory's value function. The simplified version assumes a symmetric S-shaped value function.

Finally, the results suggest that the diminishing sensitivity effect suggested by the data does not reflect risk attitude. Rather, this effect appears to reflect perceptual sensitivity: The extent to which decision makers exhibit diminishing sensitivity is highly sensitive to the presentation of the feedback. 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 Thaler and Benarzi [1995]; Camerer et al. [1997]). Under one explanation of this inconsistency, it reflects a difference between low-stakes "small decisions" (the situations examined here), and learning in high-stakes settings.

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 underinvested in the stock market, analyzed by Thaler and Benarzi [1995], can be explained through the diminishing sensitivity hypothesis supported here.[7]

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 high-stakes 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].[8] Nor do they do so in the context of riskless choice (see Shapira [1981]; Tversky and Kahneman [1991]). The current results do suggest that the loss aversion phenomenon is less general than we originally believed.

Potential implications: The example of safety rules

An attempt to derive the potential 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 high-stakes decisions. Second, the SPT 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, low-stakes decisions have consequential economic implications. Moreover, in many of these cases, the two tendencies captured by SPT 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. Indeed, the two psychological assumptions abstracted in SPT imply counterproductive risk taking in this set of situations. Diminishing sensitivity implies bias toward risk seeking because it implies insufficient sensitivity to large losses. A bias in the same direction is predicted by the assumed reliance on small samples: Since low probability events are likely to be underrepresented in small samples, the decision maker is likely to behave as if he or she believes that "it won't happen to me." Thus, the current analysis suggests that the value of the enforcement of safety rule 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. The joint effect of these tendencies can be captured with a simple model that generalizes prospect theory and probability matching.

REFERENCES

Barber, Brad M., and Terrance Odean, "Trading is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors," *Journal of Finance,* LV (2000), 773-806.

Barron, Greg, and Ido Erev, "Small Feedback-Based Decisions and Their Limited Correspondence to Description Based Decisions," *Journal of Behavioral Decision Making,* XVI (2003), 215-233.

Benartzi, Shlomo, and Richard H. Thaler, "Myopic loss aversion and the equity premium puzzle," *Quarterly Journal of Economics,* CX (1995), 73-92.

Blume, Marshall E., and Irwin Friend, "The Asset Structure of Individual Portfolios and some Implications for Utility Functions," *Journal of Finance,* XXX (1975), 585-603.

Brown, George W., "Iterative solutions of games by fictitious play," In T.C. Koopmans, editor, *Activity Analysis of Production and Allocation.* (NewYork: Wiley, 1951).

Busemeyer, Jerome R., and Julie C. Stout, "A Contribution of Cognitive Decision Models to Clinical Assessment: Decomposing Performance on the Bechara Gambling Task," *Psychological Assessment*, XIV (2002), 253-262.

Camerer, Colin F., Linda Babcock, George Loewenstein, and Richard H. Thaler, "Labor Supply of New York City Cabdrivers: One Day at a Time," *Quarterly Journal of Economics*, CXII (1997), 407-441.

Camerer, Colin F. and Teck-Hua Ho, "Violations of the Betweenness Axiom and Nonlinearity in Probability," *Journal of Risk and Uncertainty*, VIII (1994), 167-196.

Erev, Ido, and Greg Barron, ”On Adaptation, Maximization, and Reinforcement Learning Among Cognitive Strategies,” *Psychological Review*, CXII (2005), 912-931.

Erev, Ido, Ira Glozman, and Ralph Hertwig, “Context, Mere Presentation and the Impact of Rare Events,” Technion Working Paper (2006).

Erev, Ido, and Alvin E. Roth, “Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria,” *American Economic Review*, LXXXVIII (1998), 848-881.

Ert, Eyal, and Ido Erev, “The effect of losses in decisions under risk,” Technion Working Paper (2006).

Estes, William K., “Toward a statistical theory of learning,” *Psychological Review,* LVII (1950), 94-107.

Gneezy, Uri, and Jan Potters, “An experiment on risk taking and evaluation periods,” *Quarterly Journal of Economics,* CXII (1997), 631-645.

Wu, George, and Richard Gonzalez, ”Curvature of the Probability Weighting Function,” *Management Science*, XII (1996), 1676-1690.

Gray, Jeffrey A., “Three Fundamental Emotion Systems,” In P. Ekman and R. J. Davidson (Eds.), *The nature of emotion: fundamental questions* (pp. 243-247). (New York: Oxford University Press 1994).

Higgins, Tory E., “Beyond Pleasure and Pain,” *American Psychologist,* LII (1997), 1280-1300.

Kahneman, Daniel, and Amos Tversky, “Prospect Theory: An Analysis of Decision under Risk,” *Econometrica*, XLVII (1979), 263-291.

Kareev, Yaakov, “Seven (Indeed Plus Minus Two) and the Detection of Correlations,” *Psychological Review,* CVII (2000), 397-402.

Katz, Leonard, “Effects of Differential Monetary Gain and Loss on Sequential Two-Choice Behavior,” *Journal of Experimental Psychology*, LXVIII (1964), 245-249.

Kelly, Morgan, “All Their Eggs in One Basket: Portfolio Diversification of U.S. Households,” *Journal of Economic Behavior and Organization*, XXVII (1995), 87-96.

Myers, Jerome. L., and E. Sadler, “Effects of Range of Payoffs as a Variable in Risk Taking,” *Journal of Experimental Psychology*, LX (1960), 306-309.

Osborne, Martin J., and Ariel Rubinstein, “Games with Procedurally Rational Players,” *American Economic Review,* LXXXVIII (1998), 834-847.

Rabin, Matthew, “Risk Aversion and Expected Utility Theory: A Calibration Theorem,” *Econometrica,* LXVIII (2000), 1281-1292.

Ulrich Schmidt, Ulrich, and Stefan Traub, “An Experimental Test of Loss Aversion,” *Journal of Risk and Uncertainty*, XXV (2002), 233-249.

Shapira, Zur, “Making Tradeoffs Between Job Attributes,” *Organizational Behavior and Human Performance*, XXVIII (1981), 33l-355.

Thaler, Richard H., Amos Tversky, Daniel Kahneman, and Alan Schwartz, “The Effect of Myopia and Loss Aversion on Risk Taking: An Experimental Test,” *Quarterly Journal of Economics,* CXII (1997), 647-661.

Tversky, Amos, and Daniel Kahneman, ”Loss Aversion in Riskless Choice: A Reference Dependent Model,” *Quarterly Journal of Economics,* CVII (1991), 1039-1061.

Tversky, Amos, and Daniel Kahneman, ”Advances in Prospect Theory: Cumulative Representation of Uncertainty,” *Journal of Risk and Uncertainty*, V (1992), 297-323.

Wallsten, Thomas W., Thimoty J. Pleskac, and C. W. Lejuez, “Modeling a sequential risk taking task,” *Psychological Review*, CXII (2005), 862-880.

Yechiam, Eldad, Jackson Goodnight, John E. Bates, Jerome R. Busemeyer, Kenneth A. Dodge, Gregory S. Pettit, and Joseph P. Newman, “A Formal Cognitive Model of the Go/No-Go Discrimination Task: Evaluation and Implications,” *Psychological Assessment* (in press).

NOTES

1. i.e., the returns for “bonds” and “stocks” were presented in bar graphs.

2. Camerer 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 and Gonzalez [1996]).

3. This interpretation can be justified based on the arguments presented by Rabin [2000].

4. To clarify the presentation of the model we focus on capturing the behavior of a typical agent whose choice proportions equal the mean choice proportions over participants. The role of individual differences is considered below.

5. The 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.

6. The original version of RELACS includes three cognitive strategies: Slow best replay (a variant of stochastic fictitious play), Fast best replay (a variant of deterministic fictitious play with fast forgetting), and Case-based reasoning. The refined model replaces the case-based rule with the SPT rule.

7. Notice 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., Blume and Friend [1975]; Kelly [1995]; Barber and Odean [2000]).

8. However, recent research (Schmidt and Traub, [2002]; Ert and Erev [2006]) does question the robustness of the loss aversion hypothesis in decisions under risk.

TABLE I
THE PROBLEMS STUDIED IN EXPERIMENT 1

Basic problems				Proportion of S choices	
		Verbal description	Notation	Mean (SD)	Median
1	S:	0 with certainty	0	0.48	0.44
	R:	1000 with probability 0.5	(1000, .5; -1000)	(.25)	
		-1000 otherwise			
2	S:	1000 with certainty	1000	0.70	0.73
	R:	2000 with probability 0.5	(2000, .5; 0)	(.17)	
		0 otherwise			
		Mixed-Gain (MG) Score		*-0.22*	*-0.17*
				(.35)	
3	S	400 with certainty	400	0.75	0.80
	R	1400 with probability 0.5	(1400, .5; -600)	(.18)	
		-600 otherwise			
4	S	1400 with certainty	1400	0.66	0.64
	R	2400 with probability 0.5	(2400, .5; 400)	(.24)	
		400 otherwise			
		Mixed-Gain (MG) Score		*0.09*	*0.07*
				(.22)	

The left-hand columns present the 4 basic problems studied in Experiment 2. The right-hand columns present the main results over the 100 trials run in the two conditions.

TABLE II

THE PROBLEMS STUDIED IN EXPERIMENT 2

				Proportion of S choices			
Basic problems				Condition Low		Condition High	
		Verbal description	Notation	Mean (SD)	Median	Mean (SD)	Median
5	S:	A draw from the interval (0,1)	0 + u(0,1)	0.49 (.33)	0.45	0.57 (.28)	0.62
	R:	A draw from the interval (-1,0) with probability 0.5 A draw from the interval (2,3) otherwise	(-1, .5; 2) + u(0,1)				
6	S:	A draw from the interval (3,4)	3 + u(0,1)	0.55 (.25)	0.52	0.47 (.24)	0.47
	R:	A draw from the interval (2,3) with probability 0.5 A draw from the interval (5,6) otherwise	(2, .5; 5) + u(0,1)				
7	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(.25, .177, 0)	0.49 (.30)	0.47	0.60 (.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)				
8	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, .177, 12)	0.53 (.26)	0.53	0.50 (.26)	0.51
	R	A draw from a normal distribution with a mean of 13 and standard deviation of 3.54	N(13, 3.54)				
		Mixed-Gain Score (MG)		*-0.09 (.53)*	*-0.02*	*0.21 (.44)*	*0.26*

The left-hand columns present the four basic problems studied in Experiment 2. The right-hand columns present the main results over the 100 trials run in the two conditions.

TABLE III:

COMPARISON OF THE RESULTS AND THE MODEL PREDICTIONS UNDER TWO SETS OF ESTIMATED PARAMETERS

				Prop. of S choices		
Experiment	Problem	R	S	Observed	SPT est12	SPT est52
1						
0	1	(1000, .5; -1000)	0	0.48	0.50	0.50
	2	(2000, .5; 0)	1000	0.72	0.65	0.68
+400	3	(1400, .5; -600)	400	0.75	0.65	0.68
	4	(2400, .5; 400)	1400	0.65	0.58	0.59
2						
Low	5L	(2, .5; -1) +u(0,1)	0 +u(0,1)	0.49	0.46	0.37
	6L	(5, .5; 2) + u(0,1)	3+ u(0,1)	0.55	0.42	0.38
	7L	N(1.00, 3.54)	TN(0.25, 0.177) truncated at 0	0.49	0.48	0.41
	8L	N(13.00, 3.54)	TN(12.25, 0.177) truncated at 12	0.53	0.45	0.41
High	5H	(200, .5; -100) +u(0,100)	0 +u(0,100)	0.57	0.65	0.67
	6H	(500, .5; 200) + u(0,100)	200+ u(0,100)	0.47	0.44	0.44
	7H	N(100, 354)	TN(25, 17.7) truncated at 0	0.60	0.59	0.59
	8H	N(1300, 354)	TN(1225, 354) truncated at 1200	0.50	0.46	0.45

Notice that the letter added to the problem number in Experiment 2 reflects the magnitude condition: L for Low, and H for High

TABLE IV

COMPARISON OF THE RESULTS AND THE MODEL PREDICTIONS ON THE PROBLEMS STUDIED BY EREV AND BARRON (2005)

Problem and paradigm	Alternative H	Alternative L	P(L) Observed	SPT
1 MI	(11)	(10)	0.19	0.20
2 MI	(11)	(1,0.50; 19)	0.38	0.44
3 MI	(1, 0.50; 21)	(10)	0.44	0.45
4 MI	(-10)	(-11)	0.11	0.20
5 MI	(-10)	(-1, .50; -21)	0.50	0.45
6 MI	(-1, 0.50; -19)	(-11)	0.51	0.45
7 CF	(11)	(10)	0.08	0.20
8 CF	(1, 0.50; 21)	10	0.43	0.45
9 CF	-10	-11	0.07	0.20
10 CF	-10	-1, 0.50; -21	0.45	0.45
	Problems 11-14: G is the gamble (-5, .5; +5)			
11 MI	N(21,3)	N(18,3)	0.21	0.28
12 MI	N(21,3)	N(18,3) + G	0.29	0.34
13 MI	N(21,3)+G	N(18,3)	0.45	0.34
14 MI	N(21,3)+G	N(18,3) + G	0.42	0.37
	Problems 15-20: H= (x if E; -x if not-E) L=(-x if E; x if not-E)			
15 PL	x=1, p(E) =.6		0.46	0.40
16 PL	x =10, p(E) =.6		0.47	0.40
17 PL	x=1, p(E) =.7		0.38	0.31
18 PL	x =10, p(E) =.7		0.36	0.31
19 PL	x=1, p(E) =.8		0.27	0.25
20 PL	x =10, p(E) =.8		0.27	0.25
21 MI	(4, 0.80; 0)	(3)	0.39	0.38
22 MI	(4, 0.20; 0)	(3, 0.25; 0)	0.48	0.48
23 MI	(32; 0.10, 0)	(3)	0.73	0.64
24 MI	(32, 0.025; 0)	(3, 0.25; 0)	0.62	0.62
25 MI	(-3)	(-32, 0.10; 0)	0.58	0.64
			0.41	0.38
26 MI	N(100,354)	TN(25, 17.7)	0.71	0.60
27 MI	N(1300, 354)	N(1225, 17.7)	0.53	0.44
28 MI	N(1300, 17.7)	N(1225, 17.7)	0.22	0.20
	Problems 29-35: H=(G if E; B if not-E) L=(B if E; G if not-E)			
29 PL	P(E) =.75, G=5, B=0		0.22	0.28
30 PL	P(E)=.75, G=5, B= -5		0.05	0.27
31 PL	P(E)=.7, G=6, B= 2		0.32	0.31
32 PL	P(E)=.7, G=4, B= 0		0.33	0.31
33 PL	P(E)=.7, G=2, B= -2		0.23	0.31
34 PL	P(E)=.7, G=0, B= -4		0.19	0.31
35 PL	P(E)=.7, G=-2, B= -6		0.28	0.31
	Problems 36-39: P(E) =.7; P(F)=.9 H=(G if E&F, B if not-E&F, 0 otherwise) L = (B if E&F, G if not-E&F, 0 otherwise)			
36 PL	G=6, B=2		0.27	0.33
37 PL	G=-2, B=-6		0.24	0.34
38 MI	G=6, B=2		0.35	0.33
39 MI	G=-2, B=-6		0.33	0.33
40 MI	(-3)	(-4, 0.80; 0)	0.41	0.38

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 if E; y 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. Stars (*) stand for a lower maximization rate than the prediction of probability matching. The right-hand column presents the prediction of the SPT model.

FIGURE I

PROPORTION OF SAFE CHOICES IN 10 BLOCKS OF 10 TRIALS IN EACH OF THE FOUR PROBLEMS STUDIED IN EXPERIMENT 1

I.a Problems 1 & 2:

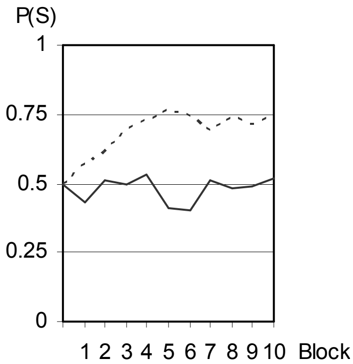

Problem		S	R
— 1	Mixed	0	(1000, .5; -1000)
···· 2	Gain	1000	(2000, .5; 0)

I.b Problems 3 & 4:

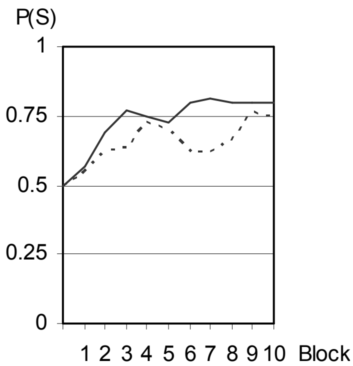

Problem		S	R
— 3	Mixed	400	(1400, .5; -600)
···· 4	Gain	1400	(2400, .5; 400)

The notation (*X, p; Y*) refers to a gamble that yields a payoff of *x* with probability *p* and *y* otherwise.

FIGURE II

PROPORTION OF SAFE CHOICES IN 10 BLOCKS OF 10 TRIALS IN EACH OF THE FOUR PROBLEMS STUDIED IN EXPERIMENT 2

II.a Problems 5 & 6:

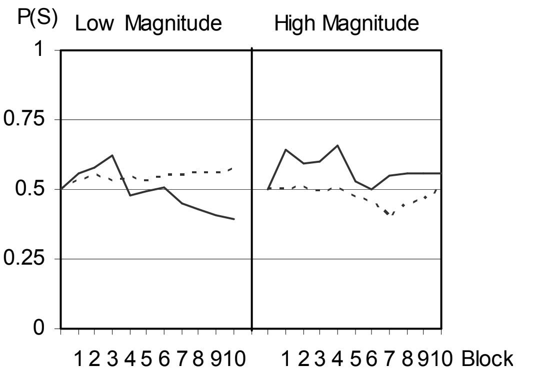

Problem		S	R
— 5	Mixed	0 + u(0,1)	(2, .5; -1)+ u(0,1)
···· 6	Gain	3 + u(0,1)	(5, .5; 2)+ u(0,1)

II.b Problems 7 & 8:

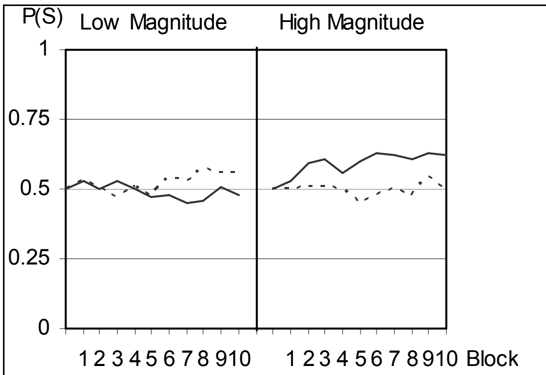

Problem		S	R
— 7	Mixed	TN(0.25, 0.177, 0)	N(1.00, 3.54)
···· 8	Gain	TN(12.25, 3.54, 12)	N(13.00, 3.54)

The left-and right-hand columns present the results in Conditions Low (point magnitude) and High respectively. 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 N(*B, F*) refers to a draw from a normal distribution with mean of *B* and standard deviation of *F*. In TN(*B, F, T*) the payoff is truncated at T.