Quantitative predictions in social science, and the choice prediction competition

Ido Erev, Eyal Ert and Alvin E. Roth

Max Wertheimer Minerva Center for Cognitive Studies, Faculty of Industrial Engineering and Management, Technion, Haifa, Israel,
and Harvard Business School, Boston, USA

June 30, 2008

Abstract:

Behavioral decision research is in a position to reduce the gap between the exact and the social sciences. The analysis of decision tasks allows quantitative prediction of behavior. The main goal of the current paper is to clarify and further this claim. We organized three open choice prediction competitions that are focused on three related choice tasks: one shot decisions under risk, one shot decisions from experience, and repeated decisions from experience. Each competition was based on two experimental studies: An estimation study, and a competition study. Both studies used the same method and subject pool, and examined randomly selected decision problems from the same distribution. After the termination of the estimation study we posted the results and their fit with several baseline models on the web (http://tx.technion.ac.il/~eyalert/Comp.html), and challenged researchers to participate in competitions that focus on the prediction of the results of the second (competition) study. Results from the estimation study reveal that some of the baseline models provide impressive predictions. The results also highlight the robustness of the difference (significant negative correlations) between decisions from description and decisions from experience. The results of the competitions will be presented in the next draft of the current paper.

Keywords: Fitting; Generalization criteria; prospect theory; reinforcement learning; naïve sampling, equivalent number of observations (ENO).

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.

The Merriam Webster dictionary defines "exact science" as **"**a science (as physics, chemistry, or astronomy) whose laws are capable of accurate quantitative expression." Most observers, including most social scientists, regard the social sciences as "inexact" by this definition, since theories in the social sciences are mostly used to make only qualitative predictions. A clear indication of the significance of the gap between the exact and the social sciences is provided by an analysis of exams used to evaluate college students. Typical questions in the exact sciences ask the examinees to predict the outcome of a particular experiment, while typical questions in the social sciences ask the examinees to exhibit understanding of a particular theoretical construct (see Erev & Livne-Tarandach, 2005 analysis of the GRE exams).

The analysis of decision tasks can reduce this gap. The quantification of theories involves two steps: An abstraction of the environment as an individual choice task (or as a multiplayer game), and an abstraction of the choice rule. Many choice theories have been given precise quantifications in this way. For example, Tversky and Kahneman (1992) present a refined variant of prospect theory (Tversky & Kahneman, 1979) that includes explicit quantification of the relevant assumptions. A focus on quantitative expressions is also common in other refinements of prospect theory (e.g., Wu & Gonzalez, 1996; Prelec 1998) and in leading alternatives to this theory (see e.g., Birnbaum & Navarrete, 1998; Brandstätter, Gigerenzer & Hertwig, 2006).

However, most previous analyses of quantitative models of choice behavior focus on fitting the results of a particular set of experiments, and stop short of evaluating the predictive value of the fitted models. That is, the models' parameters are estimated to fit the target data, and the accuracy of the fit is used to compare models. We believe that one of the reasons for the tendency to avoid evaluation of quantitative predictions is a product of a problematic incentive structure. The evaluation of quantitative predictions tends to be expensive and less interesting than the popular analyses. Research that leaves the evaluation of the quantitative predictions to future studies can start with a presentation of few interesting phenomena, and concludes with the presentation of an elegant and insightful model that captures them. To study quantitative predictions, on the other hand, researchers have to consider a wide set of randomly selected problems, and random samples of problems are less immediately appealing than specially selected cases meant to illustrate a qualitative point. The researcher then has to estimate models, and then run another large and boring (random sample) study in order to compare the different models. Moreover, the cost can increase if the researchers' favorite models fail.

The main goal of the current research is to contribute to the quantification of descriptive models of choice behavior by facilitating the transition from fitting data to quantitative predictions of behavior in a new set of choice problems. It tries to address the high cost and boredom problems by the organization of an open choice prediction competitions that can change the incentive structure. We plan to run the necessary boring studies, and challenge other researchers to predict the results.

A second goal of the current research is to clarify the relationship between decisions that are made based on a description of the payoff distributions (like the situations examined by Kahneman & Tversky, 1979), and decisions that are made based on experience. Recent research reveals large gaps between these environments (see Barron & Erev, 2003; Weber et al., 2004; Hertwig et al., 2004; Hau et al., 2008; Newell Demes & Rakow, in press; Ungemach, Chater & Stewart, 2008). Decision makers tend to overweight rare events when they have to rely on a description of these events, and tend to underweight rare events when they rely on personal experience.

1. Method

The current research involves three related, but independent, choice prediction competitions. All three competitions focus on the prediction of binary choices between a safe prospect that provides Medium payoff (referred to as M) with certainty, and a risky prospect that yields High payoff (H) with probability Ph, and Low payoff (L) otherwise. Thus, the basic choice problem is:

S: M with certainty

R: H with probability Ph; L otherwise

Table 1 presents 60 problems of this type that will be considered below. Each competition will focus on a distinct experimental condition. Three conditions (and competitions) will be considered. In Condition Description the participants are asked to make a single choice based on a description of the prospects (as in decisions under risk paradigm considered by Kahneman & Tversky, 1979). Condition Experience-sampling (E-sampling) focuses on one-shot decisions from experience (as in Hertwig et al., 2004), and Condition Experience-repeated (E-repeated) focuses on repeated decisions from experience (as in Barron & Erev, 2003).

<Insert Table 1>

The three competitions will be based on two experimental studies: An estimation study, and a competition study. The two studies use the same (three conditions) method and examine similar, but not identical, decision problems and decision makers. The estimation study was run in March 2008. After the termination of this study we posted the results (described in Table 1) on the web (see http://tx.technion.ac.il/~eyalert/Comp.html) and challenged researchers to participate in three competitions that focus on the prediction of the results of the second (competition) study. Each competition focuses on one experimental condition. The competition study was run in May 2008, but we will not look at the results until September 2nd 2008. The predictions submission deadline is September 1st 2008.

Researchers participating in the competitions are allowed to study the results of the estimation study. Their goal is to develop a model that will predict the results of the competition study. The model should be implemented in a computer program that reads the payoff distributions of the relevant gambles as an input and predicts the proportion of risky choices as an output. Thus, we use the generalization criterion methodology (see Busemeyer & Wang, 2000) and the implied competitions can be described as simplified variants of the competition organized by Arifovic, McKelvey, and Pevnitskaya (2006).

1.1 The problem selection algorithm.

Each study focused on 60 problems. The exact problems were determined with a random selection of the parameters M, H, Ph and L using the algorithm described in Appendix 1. Notice that the algorithm implies that about 1/3 of the problems involve rare (low probability) High outcomes ($Ph < .1$), and about 1/3 involve rare Low outcomes ($Ph > .9$). In addition the algorithm implies that 1/3 of problems are in the gain domain (all outcomes are positive), 1/3 are in the loss domain (all outcomes are negative), and the rest are mixed problems (at least one positive and one negative outcome).

Table 1 presents the 60 problems that were selected for the estimation study. The same algorithm was used to select the 60 problems in the competition study. Thus, the two studies focused on choice problems that were sampled from the same space of problems.

1.2 The estimation study

One hundred and sixty Technion students participated in the estimation study. Participants were paid 40 Sheqels ($11.4) for showing up, and could earn more money or lose part of the show-up fee during the experiment. Each participant was randomly assigned to one of the three experimental conditions.

Each participant was seated in front of a personal computer and was presented with a sequence of choice tasks. The exact tasks depended on the experimental condition as explained below. The procedure lasted about 40 minutes on average in all three conditions.

The payoffs on the experimental screen in all conditions referred to Israeli Sheqels. At the end of the experiment one choice was randomly selected and the participant's payoff for this choice determined his/her final payoff.

The 60 choice problems listed in Table 1 (the estimation set) were studied under all three conditions. The main difference between the three conditions was the source of the information (description, sampling or feedback). Yet, the manipulation of this factor implied other differences. The unique properties of the experimental method in each of the three conditions are described below:

Condition Description (One-shot decisions under risk):

Twenty Technion students were assigned to this condition. Each participant was seated in front of a personal computer screen and was then presented with each of the 60 problems. Participants were asked to choose once between the sure payoff and the risky gamble in each of the 60 problems that were randomly ordered. A typical screen and the instructions are presented in Appendix 2.

Condition Experience-Sampling (E-sampling, one shot decisions from experience)

Forty Technion students participated in this condition. They were randomly assigned to two different sub-groups. Each sub-group contained 20 participants who were presented with a representative sample of 30 problems from the estimation set (each problem appeared in only one of the samples, and each sample included 10 problems from each payoff domain). The participants were told that the experiment includes several games, and in each game they will be asked to choose once between two decks of cards (represented by two buttons on the screen). It was explained that before making this choice they will be able to sample the two decks. Each game was started with the sampling stage, and the participants were asked to press the "choice stage" key when they felt that they have sampled enough (but not before sampling at least once from each deck).

The outcomes of the sampling were determined by the relevant problem. One deck corresponded to the safe alternative: All the (virtual) cards in this deck provided the medium payoff. The second deck corresponded to the payoff distribution of the risky option; e.g., sampling the risky deck in problem 21 resulted with the payoff “+2 Sheqels” in 10% of the cases, and outcome “-5.7 Sheqels” in the other cases.

At the choice stage participants were asked to select once between the two virtual decks of cards. Their choice yielded a random draw of one card from the selected deck and was considered at the end of the experiment to determine the final payoff. A typical screen and the instructions are presented in Appendix 2.

Condition Experience-repeated (E-repeated, repeated decisions from experience):

One-hundred Technion students participated in this condition. They were randomly assigned to five different sub-groups. Each sub-group contained 20 participants who were presented with 12 problems (each problem appeared in only one of the samples, and each sample included equal proportion of problems from each payoff domain). Each participant was seated in front of a personal computer and was presented with each of the problems for a block of 100 trials. Participants were told that the experiment would include several independent sections (each section included a repeated play of one of the 12 problems), in each of which they would be asked to select between two unmarked buttons that appeared on the screen (one button was associated with the safe alternative and the other button corresponded to the risky gamble of the relevant problem) in each of an unspecified number of trials. Each selection followed with a presentation of its outcome in Sheqels (a draw from the distribution associated with that button, e.g., selecting the risky button in problem 21 resulted with gain of 2 Sheqels with probability 0.1 and loss of 5.7 Sheqels otherwise). Thus, the feedback was limited to the obtained payoff; the forgone payoff (the payoff from the unselected button) was not presented. A typical screen and the instructions are presented in Appendix 2.

1.3 The competition study

The competition study was identical to the estimation study with two exceptions: Different problems, and different participants. The (60) problems were selected using the algorithm used to draw the problem in the estimation studies. The (160) participants were drawn from the same population used in Study 1 (Technion Students) without replacement. That is, the participants in the competition study did not participate in the estimation study, and the choice problems were new problems randomly drawn from the same distribution

1.4 The competition criterion: Mean Squared Distance (MSD) and the Equivalent Number of Observations (ENO) interpretation.

The current competitions focus on a Mean Squared Distance (MSD) criterion. Specifically, the winner in each competition will be the model that minimizes the squared distance between the prediction and the mean observed choice proportion in the relevant condition (the mean over the 20 participants in Conditions Description and E-sampling, and over the 20 participants and 100 trials in Condition E-repeated). The main advantage of this measure is its relationship to traditional statistics (like regression, t-test and the d statistic) and intuitive interpretation.[1] These attractive features are clarified here with the

[1] Studies that focus on fitting tend to prefer the log likelihood statistic over MSD. Log likelihood main advantage is the facilitation of statistical test of the significance of the contribution of additional parameter. This advantage is not important in the current prediction tasks. It's disadvantages include a bias against (incorrect) models that make extreme predictions. For an example consider two errors: Error 1: "a

computation of the ENO (Equivalent Number of Observations) order-maintaining transformation of the MSD scores (see Erev et al., 2007). The ENO of a model is an estimation of the size of the experiment that has to be run to obtain predictions that are more accurate than the model's prediction. For example, if a model has an ENO of 10, its prediction of the probability of R choice in a particular problem is expected to be as accurate as the prediction that is based on the observed proportion of R choices in an experimental study of that problem with 10 participants. Erev et al. show that this score can be estimated as ENO = $S^2/(MSE - S^2)$ where S^2 is the pooled estimated variance over problems, and MSE is the mean squared distance between the prediction and the choices of the individual subjects (0 or 1 in the current case). When the sample size is $n = 20$, MSE = MSD + $S^2(20/19)$.

2. The results of the estimation study

The right hand columns in Table 1 present the aggregate results of the estimation study. They show the mean choice proportions of the risky prospect (the R-rate) and the mean samples that participants took in condition E-sampling. The lower panel presents the mean R-rate over problems.

2.1 Correlation analyses

Two sets of correlation analyses were conducted. The first set focuses on the relationship between the R-rates in the three conditions using problem as a unit of analysis. The results, presented Table 2a, reveal high correspondence between the two experience conditions (r[E-sampling,E-repreated] = 0.83, $p < .0001$), and a large difference between these conditions and the description condition (r[Description, E-sampling) = -0.45, $p<.05$; and r[Description, E-repeated] = -0.28 $p < .05$). This difference between the three conditions is clarified by Figure 1 that presents the R-rate as a function of Ph by condition. The results reveal an increase in the R-rate with Ph in the two experience conditions, and a decrease in the description condition. This pattern is

---

prediction 0 when the observed proportion is 1/1000." Error 2: "a prediction 0.01 when the observed proportion is 0.99." The log likelihood measure implies that a model that makes error 2 is preferred over a model that makes error 1.

consistent with the assertion that people exhibit overweighting of rare events in decisions from description, and underweighting of rare events in decisions from experience (see Barron & Erev, 2003). Additional analysis reveals that the effect of Ph holds independently of the objective attractiveness of the risky prospect.

<Insert Table 2>

<Insert Figure 1>

The second set of correlation analyses focused on individual differences. This set examines the relationship between the choices made by the same individual in different problems. Correlations of this type were computed for all the pairs of problems that were faced by the same participants. The current design allows computation of these correlations for 1770 pairs of problems in Condition Description (59+58+57…+1 = 59 * 30, when all the participants faced all 60 problems), 870 pairs in Condition E-Sampling (=29*15*2, when each of two subgroup faced 30 problems), and 330 pairs in Condition E-Repeated (=11*6*5 when each of five subgroup faced 12 problems).

The first panel in Table 2b presents the median correlations over all pairs. The results reveal positive but relatively low correlations. The median correlations are .10, .05 and .12 in Condition Description, E-sampling and E-repeated.

The second panel shows the correlation as a function of Ph (the probability of the high payoff). The results show relatively high and positive correlation between pairs of “relatively safe” problems (in which Ph is high), and between pairs of “long shot” problems (in which Ph is low). Moreover, the results also show a negative correlation between pairs of problems where one problem is relatively safe but the other is a long shot. This pattern was observed in all three conditions and can be explained with the assertion that individuals differ in a consistent fashion in their weighting of rare events. It seems that some people exhibit higher sensitivity to such events.

The third panel shows the effect of the payoff domain. The highest correlations were observed when both problems involved gains and losses. This pattern replicates previous results (see Ert& Yechiam, 2008) and can be explained with the assertion of consistent individual differences in the reaction to avoidable losses.

2.2 Learning curves

Figure 2 presents the observed R-rates in Condition E-repeated in 5 blocks of 20 trials. The 60 problems were classified to 12 graphs based on two properties: The probability of high payoff (Ph) and the relative value of the risky prospect. The most common pattern is a decrease in risky choices with experience. This pattern is predicted by the hot stove effect (Denrell & March, 2001). Comparison of the three rows suggests an interesting nonlinear relationship between the probability of high payoff (Ph) and the magnitude of the hot stove effect. A decrease in R rate with experience is clearer for high Ph and low Ph, but not for medium Ph level. This nonlinear relationship explains why previous studies that focus on gambles with equally likely outcome (like Biele et al., in press) found no evidence for the hot stove effect. The learning curves in the medium Ph problem show higher sensitivity to the expected values.

< Insert Figure 2 >

3. Baseline models

In order to clarify the challenge for the participants in the competitions we derived the predictions of several baseline models. We chose to focus on models that were described and estimated (using a Mean Squared Distance criterion) in previous research. In addition, we present two new models that were motivated by the data collected in the estimation study. The predictions of the published baseline models were derived with the original parameters, and with parameters estimated to fit the estimation set. The models are grouped by the experimental conditions that they were designed to address.

3.1 Baseline Models for Condition Description (One shot decisions under risk)

3.1.1 Original (5-parameter) Cumulative prospect theory (CPT)

According to cumulative prospect theory (Tversky & Kahneman, 1992), decision-makers are assumed to select the prospect with the highest weighted value. The weighted value of Prospect X that pays x_1 with probability p_1 and x_2 otherwise is:

(1) $$WV(X) = V(x_1)\pi(p_1) + V(x_2)\pi(p_2)$$

where $V(x_i)$ is the subjective value of outcome x_i, and $\pi(p_i)$ is the subjective weight of outcome x_i. The subjective values are given by a value function that can be described as follows:

(2) $$V(x_i) = \begin{cases} x_i^{\alpha} & \text{if} \quad x_i \geq 0 \\ -\lambda |x_i|^{\beta} & \text{if} \quad x_i < 0 \end{cases}$$

The parameters $0 < \alpha < 1$ and $0 < \beta < 1$ capture the assumption of diminishing sensitivity in the gain and the loss domain respectively. The parameter $\lambda > 1$ captures the loss aversion assertion.

The subjective weights are assumed to depend on the outcomes' rank, sign, and on a cumulative weighting function. When the two outcomes are of different sign, the weight of outcome *i* is:

(3) $$\pi(p_i) = \begin{cases} \dfrac{p_i^{\gamma}}{(p_i^{\gamma} + (1-p_i)^{\gamma})^{1/\gamma}} & \text{if} \quad x_i \geq 0 \\ \dfrac{p_i^{\delta}}{(p_i^{\delta} + (1-p_i)^{\delta})^{1/\delta}} & \text{if} \quad x_i < 0 \end{cases}$$

The parameters $0 < \gamma < 1$ and $0 < \delta < 1$ capture the tendency to overweight low-probability outcomes.

When the outcomes are of the same sign, the weight of the most extreme outcome (largest absolute value) is computed with equation (3) (as if it is the sole outcome of that sign), and the weight of the less extreme outcome is the difference between that value and 1.

The predictions of CPT were estimated with three sets of parameters: Tversky and Kahneman's parameters (for the median decision maker), the parameters estimated by Ert and Erev (2007), and the parameters that best fit the current data. The exact predictions are presented in the competition web site. Table 3a presents four measures of the accuracy of these predictions. The first two measures are the proportion of agreement between the median choice and the prediction (Pagree) and the correlation between the observed and the predicted results. These measures show high agreement (above 90%) and high correlation (above 0.84) in all cases. The third measure, and the focus of the current competition, is a Mean Square Distance (MSD) score. It reflects the mean of the squared distance of the prediction from the mean results (over participants) in each problem. Thus it is the mean of 60 squared distance scores. The final measure is the ENO transformation of the MSD score. The implied ENO scores are around 2.

< Insert Table 3 >

3.1.2 Stochastic cumulative prospect theory (SCPT).

The second model considered here is the stochastic variant of cumulative prospect theory proposed by Erev, Roth, Slonim and Barron (2002, and see a similar idea in Busemeyer, 1985). The model assume that the probability of selecting the risky prospect (R) over the safe prospect (S) is

(4) $$\Pr(R) = \frac{e^{WV(R)(\mu/D)}}{e^{WV(R)(\mu/D)} + e^{WV(S)(\mu/D)}}$$

The parameter μ captures payoff sensitivity, and D is the absolute distance between the cumulative functions implied by the two prospects. The computation of D requires a normalization of the weights of the different outcomes. The normalized weight of outcome x_1 is

(5) $$\pi_n(x_1) = \frac{\pi(x_1)}{\pi(x_1)+\pi(x_2)}$$

Assuming $x_1 > x_2$, the cumulative normal value of gamble X at point z ($0 \leq z \leq 1$) is

(6) $$CNP(X,z) = \begin{cases} V(x_1) & \text{if} \quad z \leq \pi_n(x_1) \\ V(x_2) & \text{if} \quad z > \pi_n(x_1) \end{cases}$$

Following Ert and Erev (2007) we focused on a three parameters simplification of SCPT. The simplification involves the assumption of gain loss symmetry that implies: $\lambda = 1$, $\beta = \alpha$, and $\delta = \gamma$. Table 3a presents the predictions of this model with the parameters estimated by Ert and Erev (2007), and with the parameters that best fit the current data. The results show that SCPT matches the proportion of agreement of the original model and reduces the MSD score. The ENO, with the parameters estimated by Ert and Erev, is 44.9.

Fitting all five parameters to the current data increases the estimated ENO to 80. Yet, it is not clear that the ENO interpretation of the MSD score makes sense when the same data is used to fit the parameters and estimate the model.

3.1.3 The priority heuristic

According to the priority heuristic (Brandstätter, Gigerenzer, & Hertwig, 2006), decision-makers are assumed to follow a lexicographic rule that evaluate alternatives by sequential comparison of their minimum value, their maximum value, and their respective probabilities.

The priority rule for nonnegative prospects asserts that the first comparison is done between the two minimum gains, the stopping rule is determined by a free cutoff parameter s: if the minimum gains differ by s (or more) of the maximum gain then the examination is stopped and the prospect with the better minimum gain is selected. If the difference does not pass this cutoff then the two probabilities of the minimum gains are matched. The examination is stopped if these probabilities differ by s (or more) of the probability scale. In the case that this cutoff rule is also not satisfied, the decision maker selects whichever prospect with the higher maximum gain. The same procedure applies for mixed problems and for nonpositive prospects (in such a case the word "gain" is simply replaced with the word "loss").

The prediction of the Priority Heuristic is estimated with two sets of parameters: the stopping rule parameter used by Brandstätter et al. (2006) and the stopping rule parameter that best fit the current data. The results show that this model captures the median choice as well as the CPT model: Proportion of agreement between the median choice and the prediction is above 90%. However, the mean square deviation (MSD) score is relatively high, and the implied ENO is around 2.

3.2 Baseline models for Condition E-sampling (One shot decisions from experience)

3.2.1 Primed sampler

The Primed sampler model (Erev, Glozman & Hertwig, 2008) implies a simple choice rule in condition sampling: The participants are expected to take a sample of k draws from each alternative, and select the alternative with the higher sample mean. Erev et al.'s estimation of the value of k is 5. Table 3b shows that this simple model provides good approximation of the current results. Moreover, the value $k = 5$ minimizes the MSD score. The implied ENO is 11.87.

3.2.2 Primed sampler with individual differences.

Under a natural extension of the primed sampler model the exact value of the sample size differ between participants and decisions. The current model captures this idea with the assumption that the exact sample size is uniformly drawn from the integers between 1 and k. Best fit is obtained with k =9. Table 3b shows that the added variability improves the fit.

3.3 Baseline models for condition E-Repeated (repeated decisions from experience)

3.3.1 Normalized Reinforcement Learning (NRL)

The normalized reinforcement learning model (see Erev & Barron, 2005; and a similar model in Erev, Bereby-Meyer & Roth, 1999) assumes a stochastic choice rule that is similar to the SCPT rule. Specifically, the probability of selecting the risky prospect at trial t is given by:

$$\Pr(R) = \frac{e^{WV_t(R)(\mu/D_t)}}{e^{WV_t(R)(\mu/D_t)} + e^{WV_t(S)(\mu/D_t)}} \quad (7)$$

where $WV_t(k)$ is the weighted value of action k at trial t, μ is a free payoff sensitivity parameter, and D_t is a measure of experienced payoff variability. If strategy k was selected at t, its weighted value at trial t+1 is a weighted average of $WV_t(k)$ at t, and v_t the obtained payoff at t

$$WV_{t+1}(k) = \begin{cases} (1-\omega)WV_t(k) + (\omega)v_t & if.k.was.selected.at.t \\ WV_t(k) & otherwise \end{cases} \quad (8)$$

The parameter 0 < ω < 1 captures the weight of recent outcomes. The initial value, $WV_1(k)$ is assumed to equal $A(1)$—the expected payoff from random choice (e.g., A(1) in Problem 1 is .5[18.8(.8) + 7.6(.2)] +15.5).

The payoff variability term D_t is the weighted average of the difference between the obtained payoff at trial t and t-1:

$$D_{t+1} = (\frac{t-1}{t})D_t + \frac{1}{t}ABS(v_t - v_{t-1}) \quad (9)$$

where v_0 is assumed to equal $A(1)$, and D_1 is assumed to equal μ. The predictions of this model with the parameters estimated by Erev and Barron and the parameters that best fit the data are presented in Table 3c. The results reveal a relatively large advantage of the estimated parameters.

3.3.2 Basic Reinforcement Learning

The basic reinforcement learning model, considered here, is a simplification of the NRL model. The simplification involves the assumption D_t =1. Table 3c reveals that this simplification impairs the fit.

3.3.3 RELACS

Erev and Barron (2005) present a generalization of the NRL model that assumes reinforcement learning among cognitive strategies. Since this model is rather complex, we do not present its details here. Table 3c presents its predictions for the current task with the parameters estimated by Erev and Barron, and with the parameters that best fit the current data. The results reveal that RELACS provide relatively good predictions using the original parameters, but the best fit is not very good.

3.3.4 Explorative sampler

The explorative sample model (Erev, Ert & Yechiam, in press) can be summarized with the following assumptions:

A1: Exploration and exploitation. The agents are assumed to consider two cognitive strategies: exploration and exploitation. Exploration implies a random choice. The probability of exploration is 1 in the very first trial, and (when information concerning the forgone payoffs is not available) it reduces toward an asymptote (at ε) with experience. The effect of experience on the probability of exploration depends on the expected number of trials in the experiment (*T*). Exploration diminishes quickly when *T* is small, and slowly when *T* is large (in the current study *T* = 100). This assumption is quantified as follows:

(9)
$$P(Explore_t) = \varepsilon^{\frac{t-1}{t+T^{\delta}}}$$

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

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

A3: Naïve sampling. 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.

A4: Sampling algorithm. The value of m_t at trial t is assumed to be randomly selected from the set {1, 2,…… *k*} where *k* is a free parameter.

The sampling algorithm is assumed to depend on the available information. When the 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.

A5: 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 w is a free parameter and $A_j(t)$ is the average outcome from the relevant alternative.[2]

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

$$sv(x) = \begin{cases} R_x^{\alpha_t} & if \quad R_x \geq 0 \\ -(-R_x)^{\alpha_t} & if \quad R_x < 0 \end{cases} \tag{10}$$

Where $\alpha_t = (1+V_t)^{(-\beta)}$, $\beta \geq 0$ 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 β captures the effect of diminishing sensitivity: large β 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.

Table 3c presents the predictions of the explorative sampler model with the parameters estimated by Erev et al. (in press), and with the parameters that best fit the current data. The results reveal that the model tends to over-predict the tendency to select the risky prospect.

---

[2] Implicit in this regressiveness (the assumption W > 0) is the assumption that all the experiences are weighted (because all the experiences affect the mean). The value of this implicit assumption was demonstrated by Lebiere, Gonzalez and Martin (2007).

3.3.5 Explorative sampler with recency

The last model presented here is a refinement of the explorative sampler model that was developed to capture the bias considered above. Specifically, the refined model assumes that the most recent outcome with each alternative is always considered. This change, that increased the assumed hot stove effect, is implemented by replacing assumption A3 with the following assumption:

A3’: Naïve sampling with recency. Under exploitation the agent draws (with replacement) a sample of m_t past experiences with each alternative. The first draw is the most recent experience with each alternative. All previous experiences are equally likely to be sampled in the remaining m_t-1 draws. Notice that Assumption A3’ implies a hot stove effect (see Denrell & March, 2001): An increase in risk aversion with experience.

The right hand column in Table 3c presents the predictions of the refined model. The results show that the refinement improves the fit. Additional analysis shows that the added recency effect does not impair the predictions of the explorative sampling models in the experimental conditions reviewed by Erev and Haruvy (2008).

3.4 Summary

The estimation study highlights the value of the four working assumptions that underlie the current project. First, the high equivalent number of the observations (ENO) of some of the baseline models clarifies the assertion that behavioral decision research is in a position to reduce the gap between the exact and the social sciences.

The support to the assertion that “evaluation of quantitative predictions tend to be boring and costly” includes the fact that we had to use many pages with technical details to describe the estimation study.

A third observation involves the robustness of the difference between decisions from description and decisions for experience. The estimation study reveals negative significant correlations between the two classes of choice tasks in the wide set of problems considered here.

Finally, and most importantly, the results suggest that the evaluation of predictive value can be important. One indication for this assertion is provided by the large difference of the ENO of similar models supported in previous research. For example, the addition of a stochastic response rule to CPT increases its ENO from 2 to 45.

We believe that the results of the competitions, to be reported in the next draft of the current paper, have the potential of including interesting surprises that will further clarify the assertion that the study of quantitative predictions can reduce the gap between the exact and the social sciences.

Appendix 1: The problem selection algorithm.

The 60 problems in each set are determined according to the following algorithm:

- The probability p is drawn (with equal probability) from one of the following sets (.01-.09), (.1-.9), (.91-.99)
- Two random draws are generated for the risky option (Xmax, Xmin):
- Xmin is drawn (with equal probability) from (-10, 0); Xmax is drawn from (0, +10).
- H'= Round(Xmax, .1)[3]
- L'= Round(Xmin, .1)
- The Expected Value of the risky option is determined and an error term is added to create the value of the safe option:
- m = round(H'*p+L'*(1-p), .1);
- sd = min(abs(m-L')/2,abs(m-H')/2,2); e=rannor(0)*sd; m=m+e;
- Finally the dataset is balanced to include equal proportion of problems that include nonpositive payoffs (loss domain), nonnegative payoffs (gain domain) and both positive and negative payoffs (mixed domain).
- If problem <21 then con = -max+min;
- If 20 < problem < 41 then con = 0;
- If problem > 40 then con = +max-min;
- L = L'+con; M = round(m+con,.1); H = H'+con;

[3] The function Round(x, .1) implies rounding x to the nearest decimal. The function abs(x) returns the absolute value of x. the function rannor(0) returns a randomly selected value from a normal distribution with a mean of 0 and standard deviation of 1.

Appendix 2: Translation of the instructions and typical experimental screens of each of the three conditions (Description, Experience-Sampling, and Experience-Repeated).

Condition Description:

This experiment includes several games. In each game you will be asked to select one of two alternatives.

At the end of the experiment one of the games will be randomly drawn (all the games are equally likely to be drawn), and the alternative selected in this game will be realized.

Your payoff for the experiment will be the outcome (in Sheqels) of this game.

Good luck!

Experimental Screen after selecting the safer option in Problem 32:

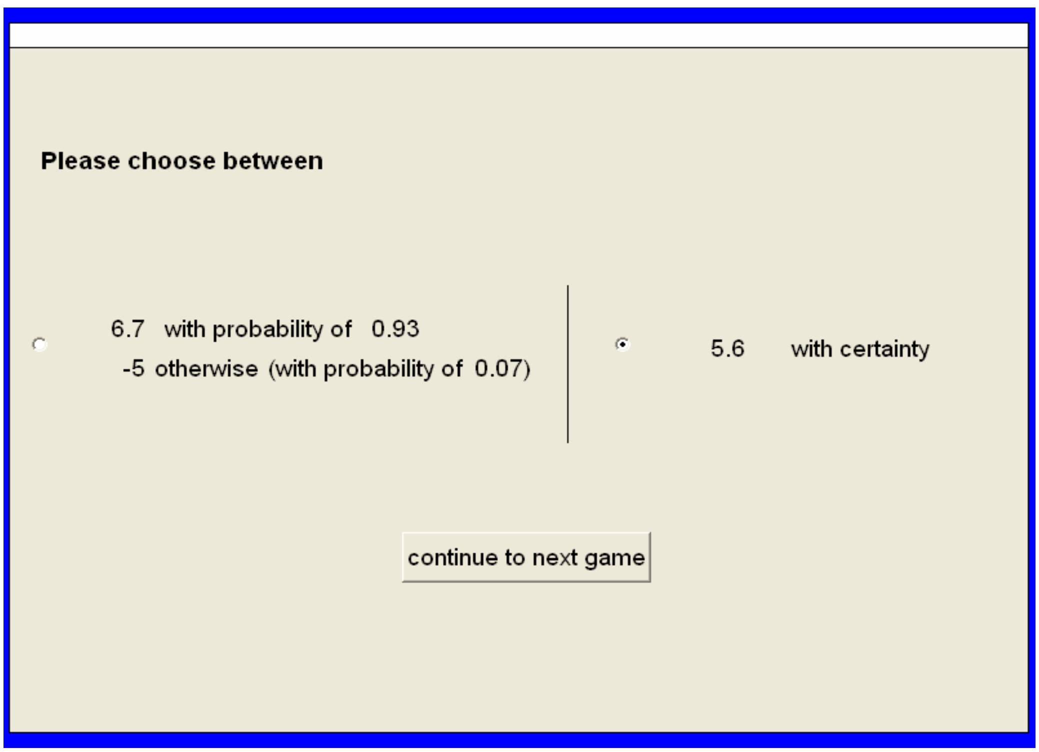

Condition E-Sampling:

This experiment includes several games. Each game includes two stages: The sampling stage and the choice stage.

At the choice stage (the second stage) you will be asked to select once between two virtual decks cards (two buttons). Your choice will lead to a random draw of one card from this deck, and the number written on the card will be the "game's outcome." During the sampling stage (the first stage) you will be able to sample the two decks. When you feel that you have sampled enough press the "choice stage" key to move to the choice stage.

At the end of the experiment one of the games will be randomly drawn (all the games are equally likely to be drawn). Your payoff for the experiment will be the outcome (in Sheqels) of this game.

Good luck!

Experimental screen (a) after sampling the deck associated with the safer option in Problem 4 during the sampling stage:

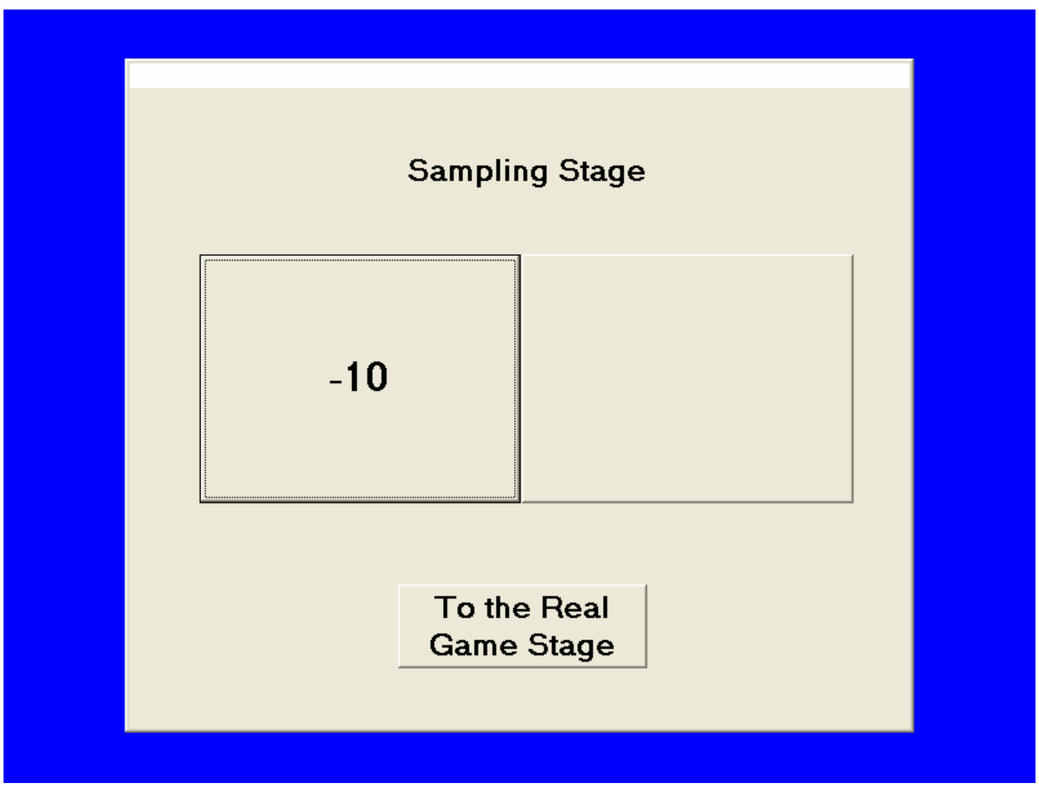

Experimental screen (b) – After choosing the deck associated with the safer option in Problem 4 during the real game stage:

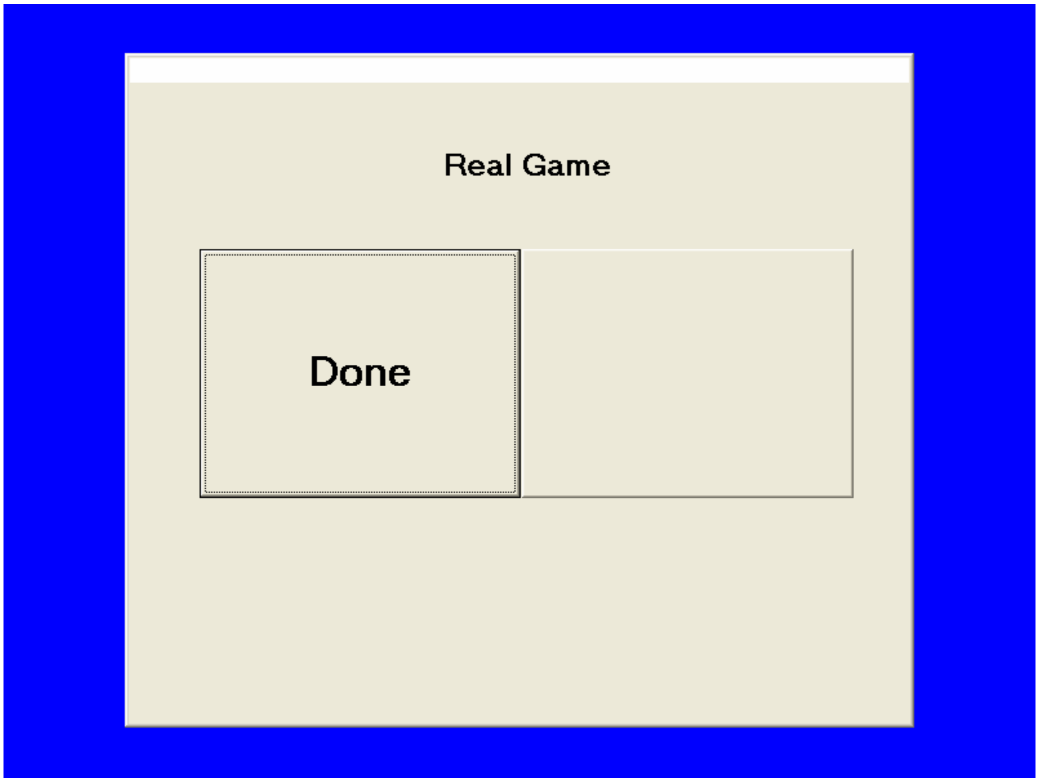

Condition E-Repeated:

This experiment includes several games. Each game includes several trials. You will receive a message before the beginning of each game.

In each trial you will be asked to select one of two buttons. Each press will result with a payoff that will be presented on the selected button.

At the end of the experiment one of the trials will be randomly drawn (all the trials are equally likely to be drawn). Your payoff for the experiment will be the outcome (in Sheqels) of this trial.

Good luck!

Experimental Screen after choosing the risky alternative in Problem 36:

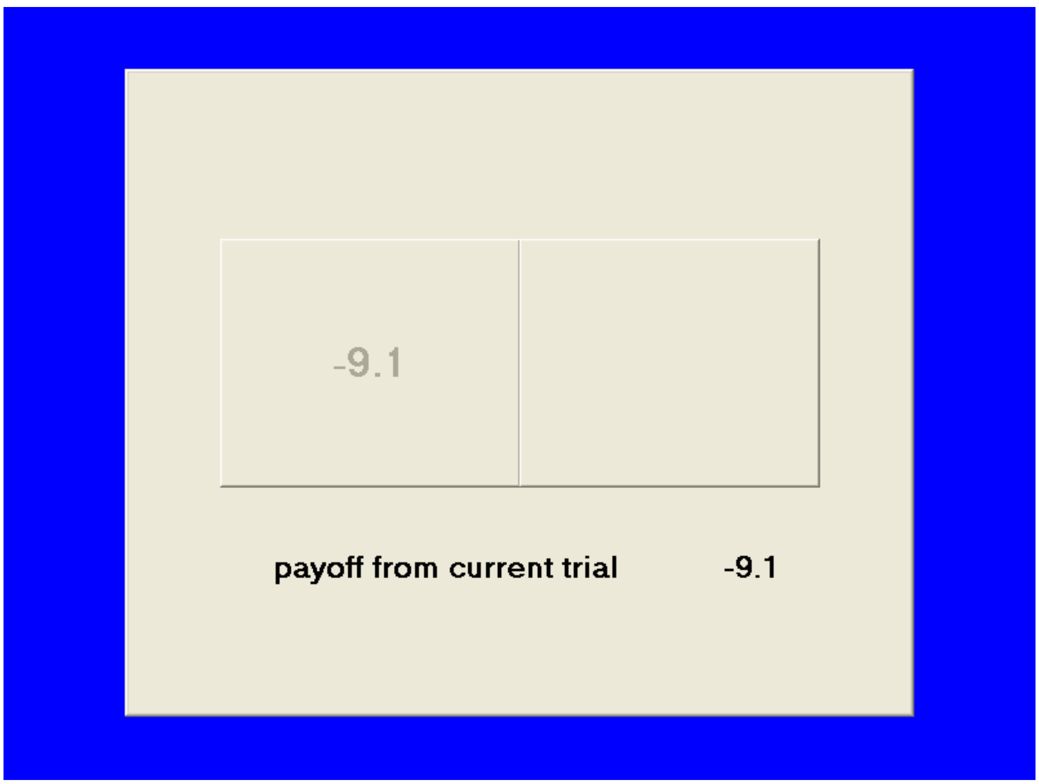

References

Arifovic, J., McKelvey, R. D., & Pevnitskaya, S. (2006). An initial implementation of the Turing tournament to learning in repeated two-person games. *Games and Economic Behavior*, *57*, 93-122.

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.

Biele, G., Erev I., & Ert, E. (in press). Learning, risk attitude and hot stoves in restless bandit problems. *Journal of Mathematical Psychology*.

Birnbaum, M. H. & Navarrete J. B. (1998). Testing descriptive utility theories: Violations of stochastic dominance and cumulative independence, *Journal of Risk and Uncertainty*, *17*, 49-79.

Brandstätter E., Gigerenzer, G., & Hertwig, R. (2006). The priority heuristic: Making choices without tradeoffs. *Psychological Review*, *113*, 409-432.

Busemeyer, J. R., (1985). Decision making under uncertainty: A comparison of simple scalability, fixed sample, and sequential sampling models. *Journal of Experimental Psychology*, *11*, 538-564.

Busemeyer, J. R., & Wang, Y. M. (2000). Model comparisons and model selections based on the generalization criterion methodology. *Journal of Mathematical Psychology*, *44*, 171-189.

Camerer, C. F & Ho T. H. (1999). Experience-weighted attraction learning in normal form games. *Econometrica, 67*, 827–874.

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., Ert, E., & Yechiam, E. (in press). Loss aversion, diminishing sensitivity, and the effect of experience on repeated decisions. *Journal of Behavioral Decision Making.*

Erev, I., Glozman, I., & Hertwig, R. (2008). Context, mere presentation and the impact of rare events. *Journal of Risk and Uncertainty*. *36*, 153-177.

Erev, I., & Haruvy, E., (2008). Learning and the economics of small decisions. forthcoming in J. H. Kagel & A. E. Roth (Eds.) *The Handbook of Experimental Economics*.http://www.utdallas.edu/~eeh017200/papers/LearningChapter.pdf

Erev, I., & Livne-Tarandach, R. (2005). Experiment-based exams and the difference between the behavioral and the natural sciences. In R. Zwick & A. Rapoport (Eds.), *Experimental Business Research, Vol 3*. Dordrecht, The Netherlands: Springer.

Erev, I., Roth, A. E., Slonim, R. L., & Barron, G. (2005). *Descriptive models as generators of prior beliefs with known equivalent number of observations (ENO).* Technion, Working Paper.

-------- (2007). Learning and equilibrium as useful approximations: Accuracy of prediction on randomly selected constant sum games. *Economic Theory*. *33*, 29-51.

Ert, E., & Erev, I**.** (2007). *Loss aversion in decisions under risk and the value of a symmetric simplification of prospect theory*. Technion, Working Paper. http://ssrn.com/abstract=1012022.

Ert, E., & Yechiam, E. (2008). *Are people consistent in their sensitivity to losses?* Technion, Working paper.

Hau, R., Pleskac, T. J., Kiefer, J., & Hertwig, R. (2008) The description-experience gap in risky choice: the role of sample size and experienced probabilities. *Journal of Behavioral Decision Making.*

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*, 534-39.

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

Lebiere, C., Gonzalez, C., & Martin, M. (2007). Instance-based decision making model of repeated binary choice. *Proceedings of the 8th International Conference on Cognitive Modeling*. Ann Arbor, Michigan, USA

Prelec, D. (1998). The probability weighting function. *Econometrica, 66*, 497-527.

Newell, B.R., Demes, K. A. & Rakow, T. (in press). Biased samples not mode of presentation: Re-examining the apparent underweighting of rare events in experience-based choice. *Organizational Behavior and Human Decision Processes*.

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*, 297-323.

Ungemach, C., Chater, N., & Stewart, N. (2008). *Decisions from Experience without Sampling Error*.

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.

Table 1: The 60 estimation set problems and the aggregate proportion of choices in risk in each of the experimental conditions.

	Risk			Safe	Proportion of choices in Risk (R - rate)			average number
Problem	High	P(High)	Low	Medium	Description	E-Sampling	E-repeated	of samples
1	-0.3	0.96	-2.1	-0.3	0.2	0.25	0.33	10.4
2	-0.9	0.95	-4.2	-1	0.2	0.55	0.50	9.7
3	-6.3	0.3	-15.2	-12.2	0.6	0.5	0.24	13.9
4	-10	0.2	-29.2	-25.6	0.85	0.3	0.32	10.7
5	-1.7	0.9	-3.9	-1.9	0.3	0.8	0.45	9.9
6	-6.3	0.99	-15.7	-6.4	0.35	0.75	0.68	9.9
7	-5.6	0.7	-20.2	-11.7	0.5	0.6	0.37	11.1
8	-0.7	0.1	-6.5	-6	0.75	0.2	0.27	13.9
9	-5.7	0.95	-16.3	-6.1	0.3	0.6	0.43	11.0
10	-1.5	0.92	-6.4	-1.8	0.15	0.9	0.44	11.8
11	-1.2	0.02	-12.3	-12.1	0.9	0.15	0.26	11.9
12	-5.4	0.94	-16.8	-6.4	0.1	0.65	0.55	11.2
13	-2	0.05	-10.4	-9.4	0.5	0.2	0.11	10.4
14	-8.8	0.6	-19.5	-15.5	0.7	0.8	0.66	12.1
15	-8.9	0.08	-26.3	-25.4	0.6	0.3	0.19	11.6
16	-7.1	0.07	-19.6	-18.7	0.55	0.25	0.34	11.0
17	-9.7	0.1	-24.7	-23.8	0.9	0.55	0.37	15.1
18	-4	0.2	-9.3	-8.1	0.65	0.4	0.34	11.2
19	-6.5	0.9	-17.5	-8.4	0.55	0.8	0.49	14.9
20	-4.3	0.6	-16.1	-4.5	0.05	0.2	0.08	10.9
21	2	0.1	-5.7	-4.6	0.65	0.2	0.11	8.8
22	9.6	0.91	-6.4	8.7	0.05	0.7	0.41	9.2
23	7.3	0.8	-3.6	5.6	0.15	0.7	0.39	10.7
24	9.2	0.05	-9.5	-7.5	0.5	0.05	0.08	14.6
25	7.4	0.02	-6.6	-6.4	0.9	0.1	0.19	8.9
26	6.4	0.05	-5.3	-4.9	0.65	0.15	0.20	13.4
27	1.6	0.93	-8.3	1.2	0.15	0.7	0.50	8.9
28	5.9	0.8	-0.8	4.6	0.35	0.65	0.58	10.6
29	7.9	0.92	-2.3	7	0.4	0.65	0.51	10.6
30	3	0.91	-7.7	1.4	0.4	0.7	0.41	10.0
31	6.7	0.95	-1.8	6.4	0.1	0.7	0.52	11.0
32	6.7	0.93	-5	5.6	0.25	0.55	0.49	11.0
33	7.3	0.96	-8.5	6.8	0.15	0.75	0.65	11.1
34	1.3	0.05	-4.3	-4.1	0.75	0.1	0.3	11.4
35	3	0.93	-7.2	2.2	0.25	0.55	0.44	12.8
36	5	0.08	-9.1	-7.9	0.4	0.2	0.09	14.6
37	2.1	0.8	-8.4	1.3	0.1	0.35	0.28	10.9
38	6.7	0.07	-6.2	-5.1	0.65	0.2	0.29	10.9
39	7.4	0.3	-8.2	-6.9	0.85	0.7	0.58	12.7
40	6	0.98	-1.3	5.9	0.1	0.7	0.61	13.5

41	18.8	0.8	7.6	15.5	0.35	0.6	0.52	9.0
42	17.9	0.92	7.2	17.1	0.15	0.8	0.48	10.8
43	22.9	0.06	9.6	9.2	0.75	0.9	0.88	9.9
44	10	0.96	1.7	9.9	0.2	0.7	0.56	10.1
45	2.8	0.8	1	2.2	0.55	0.7	0.48	19.4
46	17.1	0.1	6.9	8	0.45	0.2	0.32	9.2
47	24.3	0.04	9.7	10.6	0.65	0.2	0.25	11.8
48	18.2	0.98	6.9	18.1	0.1	0.75	0.59	9.0
49	13.4	0.5	3.8	9.9	0.05	0.45	0.13	8.9
50	5.8	0.04	2.7	2.8	0.7	0.2	0.35	10.0
51	13.1	0.94	3.8	12.8	0.15	0.65	0.52	9.0
52	3.5	0.09	0.1	0.5	0.35	0.25	0.26	11.9
53	25.7	0.1	8.1	11.5	0.4	0.25	0.11	9.0
54	16.5	0.01	6.9	7	0.85	0.25	0.18	13.4
55	11.4	0.97	1.9	11	0.15	0.7	0.66	9.6
56	26.5	0.94	8.3	25.2	0.2	0.5	0.53	14.3
57	11.5	0.6	3.7	7.9	0.35	0.45	0.45	10.0
58	20.8	0.99	8.9	20.7	0.25	0.65	0.63	12.9
59	10.1	0.3	4.2	6	0.45	0.45	0.32	10.1
60	8	0.92	0.8	7.7	0.2	0.55	0.44	10.2

Note - All problems involves binary choice between a sure payoff (Medium) and a risky option with two possible outcomes (High and Low). For example, Problem 1 describes a choice between loss of NIS 0.3 for sure, and a gamble that yields a loss of NIS 0.3 with probability of 0.96 and a loss of NIS 2.1 otherwise.

Table 2: Correlation analyses (Estimation study):

2a: The correlations between the R-rates in the different conditions using problem as a unit of analysis. All three correlations are significantly different from 0 ($p < .05$).

	E-sampling	E-repeated
Description	-.45	-.28
E-sampling		.83

2b. The median correlations between R-rates in different pairs of problems using participant as a unit of analysis. The numbers in parentheses present the number of pairs that were used to compute each median.

Dimension	Category	Description	E-sampling	E-repeated
Overall		.10 (1770)	.05 (870)	.12 (330)
Phigh	p≥.9 in both problems	.19 (300)	.18 (144)	.42 (56)
	p≥.9 in one problem, and ≤.1 in the second	-.06 (500)	-.11 (249)	-.22 (95)
	P≤.1 in both problems	.24 (190)	.14 (91)	.38 (38)
	Other	.10 (780)	.05 (386)	.14 (141)
Payoff domain	Gain in both problems	.10 (190)	.04 (90)	.00 (30)
	Gain in one, loss in the second	.07 (400)	.02 (200)	.07 (80)
	Loss in both	.11 (190)	.05 (90)	.03 (30)
	Mixed and Med<0 in both	.28 (66)	.19 (30)	.48 (10)
	Mixed and Med >0 in both	.18 (28)	.27 (12)	.07 (4)
	Other	.09 (896)	.05 (448)	.19 (176)
Higher EV	R, R	.10 (300)	.05 (144)	.09 (54)
	R, S	.10 (875)	.05 (437)	.09 (167)
	S, S	.06 (595)	.05 (289)	.22 (109)

Table 3: Summary of the descriptive value of the proposed baseline models in each condition of the estimation study. Pagree is the proportion of agreement between modal predictyion and the modal choice, Corr is the person correlation, MSD is mean squared deviation, ENO is the equivalent number of observations.

3a: Condition Description:

Model	Parameters	Pagree	Corr	MSD	ENO (S^2=.1860)
Cumulative Prospect Theory	Original α=.88, β=.88 λ=2.25 γ=.61, δ=.69	91%	0.84	0.099	2.05
	Symmetric α=β=.86 λ=1, γ= δ=.5	91%	0.83	0.1049	1.68
	Fitted α = β =.70 λ=1, γ = δ =.65	95%	0.85	.0932	2.22
Stochastic Cumulative Prospect Theory	Symmetric α=β=.77 (λ=1) γ=δ=.71, μ=2.04	93%	.91	.0134	44.9
	Fitted α=.89, β=.98 λ=1.5, γ= δ=.7 μ=2.15	91%	.92	.0116	80.6
Priority Rule	s=.1	95%	.81	.0999	2.05

3b: Condition E-sampling.

Model	Parameter	pagree	Corr	MSD	ENO (S^2=.1927)
Primed sampler	κ = 5	90%	0.81	.027	11.87
Primed sampler fitted with variability	κ = 9	93%	0.88	.017	29.3

3c: Condition E-repeated:

Model	Parameters	pagree	Corr	MSD est	ENO est (S^2=.0875)
Normalized Reinforcement Learning (NRL)	Original w=.01, λ=3	.74	.81	0.0236	4.55
	Fitted w=.15, λ=1.1	.76	.83	0.0092	18.13
Basic Reinforcement Learning (RL)	Original w=.03, λ=1	.53	.60	0.0301	3.40
	Fitted W=.15, λ=1	.56	.67	0.0224	4.85
RL among cog. Strategies (RELACS)	Original λ = 8 α =.0012 β =0.2, k=4	.77	.87	0.0121	11.32
	Fitted λ = 8, α =.005 β =0.1, k=4	.77	.88	0.0110	13.20
Explorative sampler (δ = 0.55, and w=.3 in all cases)	Original β=.15, ε=.08 k = 5	.73	.84	0.0189	6.02
	Fitted β=.05 ε=.12, k=20	.75	.85	0.0115	12.28
	With recency β=.10 ε =.12, k=8	.82	.88	0.0075	27.99

Figure 1: R-rate (proportion of risk taking) as a function of Ph (the probability of getting the high outcome from the risky gamble) in each of the three experimental conditions in the estimation study.

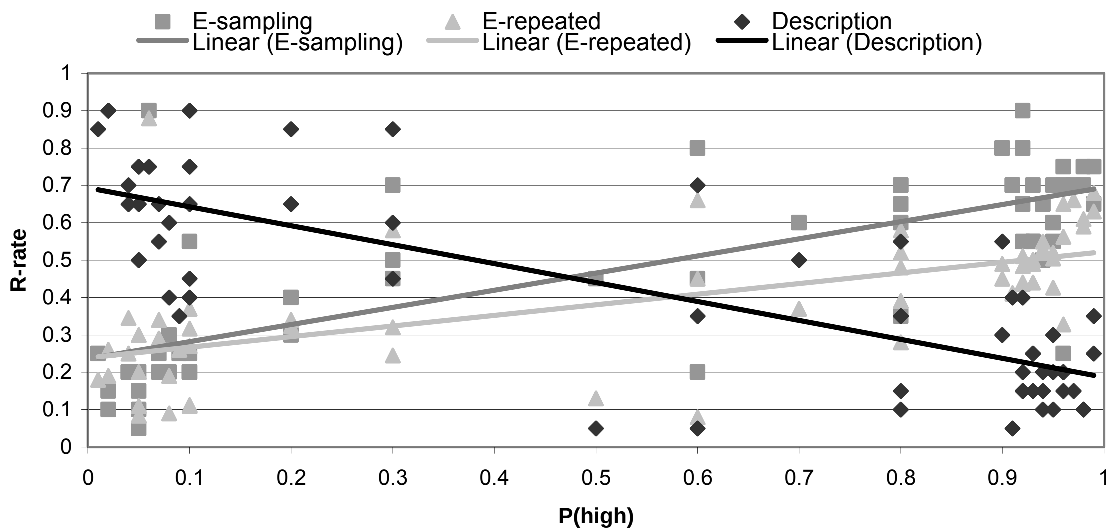

Figure 2 – observed R-rates in Condition E-repeated in 5 blocks of 20 trials. The 60 problems were classified to 15 graphs according to (a) The probability of high payoff (Ph) and the relative expected value of the risky prospect (EVr-EVs).

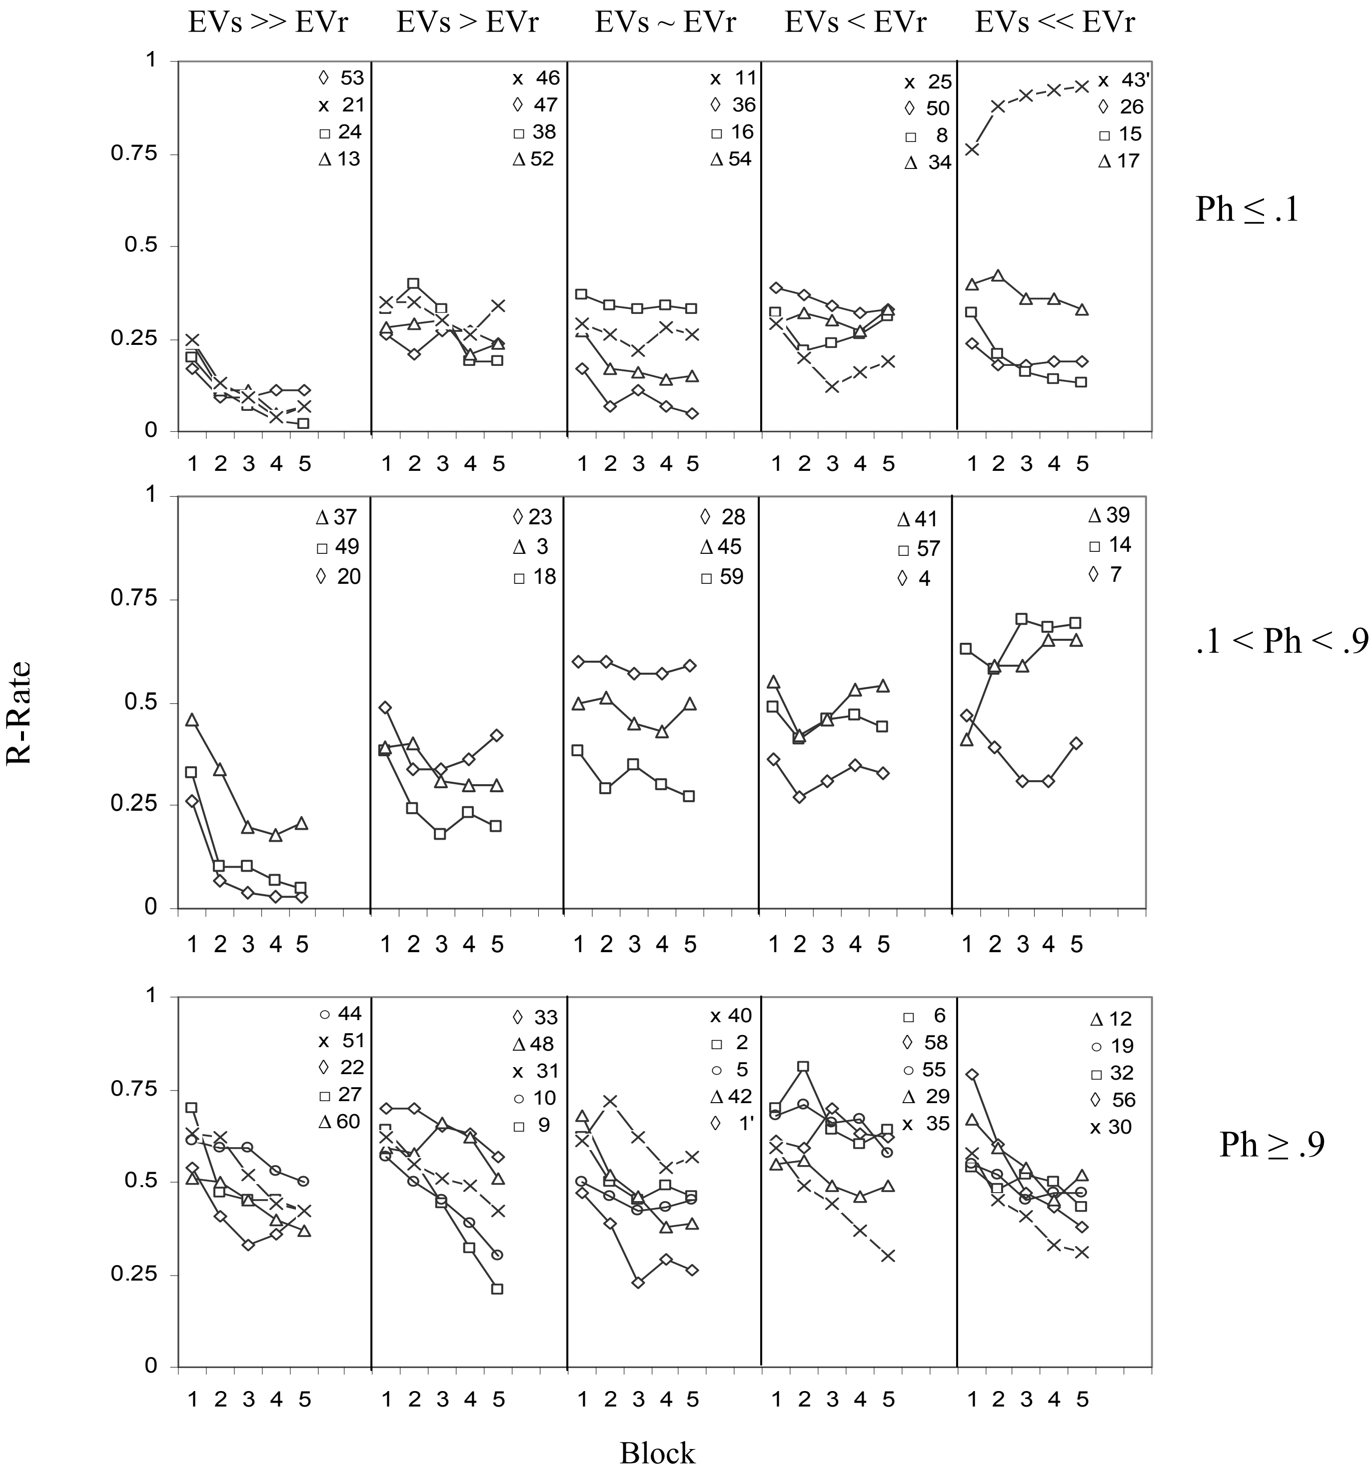

NOTE – The numbers in the legend are the problem id. In the marked problems (1 and 43) one alternative dominates the other.