

Replicated Alternatives and the Role of Confusion, Chasing, and Regret in Decisions from Experience

EYAL ERT and IDO EREV*

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

ABSTRACT

The current paper explores choice among alternatives that can be classified into distinct classes. All the members of a particular class were ‘replicated alternatives’: they promised the same payoff distribution. Information to decision makers was limited to feedback concerning the realized (obtained and foregone) payoffs. Experiment 1 demonstrates that increasing the number of replicated alternatives creates confusion (which facilitates random choice) and changes the implications of the tendency to chase recent returns (i.e., select the alternative with the best recent outcomes). This effect, termed ‘confused chasing,’ facilitates risk seeking even when this behavior impairs expected earnings. Experiment 2 reveals that increasing the number of replicated alternatives can reduce (but does not eliminate) the tendency to underweight rare events. Experiment 3 shows that the relative importance of chasing and confusion is sensitive to the likelihood of realizing lower payoffs than the forgone payoffs. The main results are summarized with a simple model assuming that payoff sensitivity decreases with experienced regret. Copyright © 2007 John Wiley & Sons, Ltd.

KEY WORDS reinforcement learning; fictitious play; momentum trading; the reversed certainty effect; independence of irrelevant alternatives; decisions from experience

INTRODUCTION

Many natural decision problems involve choices among multiple alternatives that can be classified into several classes. For example, an investment decision involves a choice between many feasible portfolios that can be classified, based on the expected returns and risks, into a small set of distinct investment policies. Similarly, gambling in a casino involves choices between many tables that can be classified based on the game they offer.

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

When the alternatives within each class are associated with similar payoff distributions, it is natural to assume that the important decision is the choice among the different classes. Indeed, mainstream research tends to simplify situations of this type in just this way. An important example involves the influential generalization of the basic study of reinforcement schedules (see Ferster & Skinner, 1957). While the basic research focuses on the choice between two possible actions (pressing or not pressing a lever), the results are generalized (e.g., Hellriegel, Slocum, & Woodman, 1998) to address behavior at the workplace. In this case, employee behavior is abstracted as a choice between two classes of actions: 'high effort' and 'low effort.' Another example involves the attempt to use basic decision-making research to capture the effect of rare terrorist attacks. Yechiam, Barron, and Erev (2005) address this problem by focusing on binary choice tasks in which decision makers are asked to select between a risky and a safe prospect.

Unspoken in these abstractions is the assumption that the implicit choice among classes (the probability that the selected alternative belongs to a particular class) is independent of the number of similar alternatives within each class. The current investigation examines the value of this assumption in a simplified environment that satisfies four conditions. First, all the alternatives within each class are assumed to have the same payoff distributions—that is, they are all 'replications' of the class's 'basic gamble.' For example, all the alternatives that belong to the class 'basic gamble H' in Problem 1 (studied in Experiment 1 below) yield a draw from a normal distribution with a mean of 11 and standard deviation of 3. Second, the realized payoffs (the draws from the basic payoff distributions) of the different alternatives are assumed to be independent. The third condition involves symmetry. The current analysis focuses on conditions in which each of the basic gambles has the same number of replications. Finally, the current analysis considers pure decisions from experience: choice situations in which decision makers are given no prior information concerning the payoff distributions of the different alternatives (and are not informed that they belong to distinct classes), but face the same problem repeatedly and observe the payoff from each alternative after each trial.

The assumption that the implicit choice between classes is independent of the number of replicated alternatives is a special case of the independence of irrelevant alternative assumptions introduced by Arrow (1959).¹ Direct examinations of the descriptive power of this assumption have documented robust violations (see Elrod, Louviere, & Davey, 1992; Goldstein & Mitzel, 1992; Tversky, 1972; Tversky & Simonson, 1993). A particularly important set of violations was observed in environments in which one class has more replicated alternatives than the other. In this environment, behavior appears to reflect a bias toward the class with more members (Benartzi & Thaler, 2001; but see also boundary conditions in Huberman & Jiang, 2006).

Analysis of the simplified set of decision situations considered here reveals two reasonable predictions concerning the effect of an increase in the number of replicated alternatives in each class. The first prediction is derived from studies of cognitive limitations (e.g., Miller, 1956). Once these limits are surpassed, behavior tends to become confused and dysfunctional (Jacoby, Speller, & Kohn, 1974; Keller & Staelin, 1987). Under this 'confusion hypothesis,' an increase in the number of replicated alternatives is likely to impair learning. Thus, it is expected to move behavior toward random choice.

A second prediction is suggested by previous studies of decisions from experience (see Barron & Erev, 2003) and investment decisions (see Hong, Lim, & Stein, 2000; Hong & Stein, 1999; Jagadeesh & Titman, 1993; Rouwenhorst, 1998). This research reveals a tendency to select the alternative with the best recent payoff. The effect of this 'chasing' of recent outcomes interacts with the number of alternatives in each class. As the number of alternatives within each class grows, the probability that one of the riskier options will be associated with the best recent outcome tends to increase accordingly. Thus, under a 'chasing hypothesis' a

¹Written as axiom C4. See also an earlier version of this condition by Nash (1950), presented as assumption No. 5. Luce (1959) presents an elegant probabilistic generalization of this principle. The basic assumption can be described as follows: if alternative x is chosen from choice set $\{w, x, y, z\}$ containing alternatives w, x, y, z , then x is considered to be at least weakly preferred to w, y, z . The value of a choice set (and subset) can, accordingly, be equated with the best alternatives available.

rise in the number of replicated alternatives is expected to increase risk seeking (see Grosskopf, Erev, & Yechiam, 2006; Shavit, Ben Zion, Erev, & Haruvy, 2005).

The three experiments described below were designed to evaluate the confusion and chasing hypotheses. The results reveal an interesting interaction between the two hypothesized effects. Experiment 1 shows that a rise in the number of replicated alternatives increases the attractiveness of the riskier class only when this class is associated with a lower expected return. These results are consistent with a joint effect of confusion and chasing. Experiments 2 and 3 shed light on the properties of this joint effect. Experiment 2 suggests that the effects of confusion and chasing are not additive. Experiment 3 shows that the relative importance of the two factors is sensitive to experienced regret. Confusion appears to drive behavior when experienced regret is likely, and chasing is more important when experienced regret is unlikely. This pattern can be summarized with a simple model assuming that payoff sensitivity decreases with experienced regret.

EXPERIMENT 1—CONFUSION VERSUS CHASING

Previous studies of decision making from experience (see review in Erev & Barron, 2005) reveal relatively flat learning curves. That is, most of the aggregate effect of experience occurs during the first 50 trials. This observation led us to use a within-participant experimental design in which each participant is faced with each of the decision problems for 50 trials. In each trial, participants were asked to select one of the available alternatives. The different alternatives included m replications of two basic gambles: Gamble H and Gamble L. The names of the two basic gambles reflect the fact that H had higher expected value (EV). Thus, each ‘class’ included m replications of one of these basic gambles. Experiment 1 focused on three variants of the following two basic choice problems:

Problem 1:

- H (EV = 11) A draw from a normal distribution with a mean of 11 and standard deviation of 3.
- L (EV = 10) A draw from a normal distribution with a mean of 10 and standard deviation of 1.

Problem 2:

- H (EV = 11) A draw from a normal distribution with a mean of 11 and standard deviation of 1.
- L (EV = 10) A draw from a normal distribution with a mean of 10 and standard deviation of 3.

Three ‘replication’ (value of m) Levels, 1, 3, and 25 were compared. In Level 1 (‘binary choice’), each basic gamble was presented once. In Levels 3 (‘6 alternatives’) and 25 (‘50 alternatives’), each basic gamble was replicated 3 and 25 times, respectively.

Each of the 50 trials under each condition opened with the presentation of $2m$ unmarked buttons on the computer screen. The m replications of each basic gamble were randomly placed before the first trial (and their location was not changed during the 50 trials). After each choice the participant was presented with both the obtained payoff and the forgone payoff (the payoff from each of the other $2m-1$ alternatives). Figure 1 presents an example of this feedback screen under each of the three m levels.

According to the ‘confusion’ hypothesis, maximization rate (selection of the class associated with Gamble H) is expected to decrease with m .² This effect is not expected to depend on the problem type. Highest maximization rate is predicted in the binary choice task ($m = 1$), and it should move toward random choice as the number of alternatives increases.

²Notice that we focus on maximization of expected value as a benchmark. We do not argue that deviations from expected value maximization reflect deviations from rational behavior. Indeed, in decisions from experience any behavior can be rational under certain prior beliefs (see Budescu, Erev, Wallsten, & Yates, 1997 for a discussion of the value of ‘rationality free’ decision research).

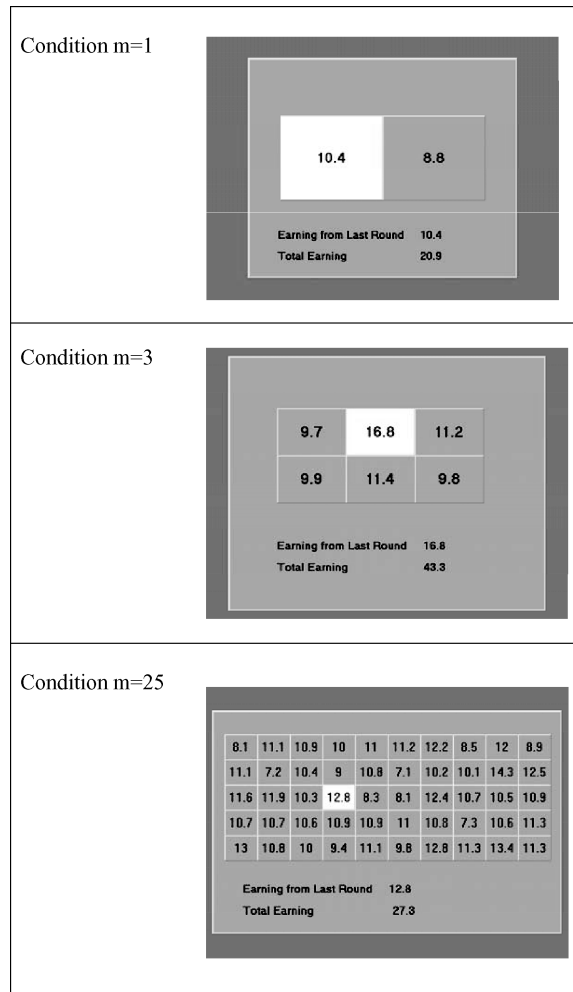


Figure 1. Typical experimental screens after the first choice in Experiment 1 in the three m levels. The selected button is highlighted (with white)

The ‘chasing’ hypothesis predicts a difference between the two problems: in Problem 1, Gamble H, which maximizes expected payoff, is associated with higher variance (i.e., is the riskier gamble). As a result, a rise in m (the number of replicated alternatives) increases the probability that chasing implies maximization. To clarify this point, the current analysis focuses on ‘perfect chasing’: a choice of the alternative with the highest most recent payoff. The predicted maximization rates under perfect chasing in Problem 1 are 64%, 95%, and 99.9% for m levels of 1, 3, and 25, respectively. Thus, under the chasing hypothesis the maximization rate in Problem 1 is expected to increase as the number of alternatives grows. The opposite pattern is predicted in Problem 2. Here, Gamble H is associated with lower variance. As a result, a rise in m decreases the probability that perfect chasing implies maximization. The exact rates are 63%, 27%, and 2% for m levels of 1, 3, and 25, respectively. Thus, under the chasing hypothesis, the maximization rate in Problem 2 is expected to fall as the number of alternatives grows.

Design and procedure

Participants were told that the experiment would include several independent sections, in each of which they would operate a different ‘computerized money machine’ with a certain number of unmarked buttons for an unspecified number of trials. In each trial, the participants were asked to select one of the buttons. After a button was pressed, the outcome associated with each button appeared on the computer screen. The participants were informed that their payoff for that trial would be the outcome that appeared on the selected button, and this value was added to their accumulated earnings. 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 conversion rate of 1 agora (about 0.023 US cent) per point.

The participants received no prior information about the game’s payoff structure, nor were they told in advance that the experiment included six sections of 50 trials each (though they were informed when a section ended and a new section was about to start). Thus, the participants had to rely on their obtained feedback: the realized payoffs after each choice.

The six sections corresponded to the six conditions (2 basic problems \times 3*m* levels). The order of the six conditions was randomized over participants. In each section, each button related to one of the alternatives; the payoffs associated with this button (alternative) were determined by random draws from the corresponding payoff distribution. For example, a selection of each of the *m* replications of the basic gamble H in Problem 1 resulted in a draw from a normal distribution with a mean of 11 and standard deviation of 1. The exact payoffs were rounded to the nearest decimal. The assignment of alternatives to buttons was randomly determined for each participant at the beginning of each section and was fixed during the section.

Participants

Twenty Technion students served as paid participants in the experiment. Payment was made according to participants’ obtained payoffs in the experimental task. Final payoffs ranged between 31 sheqels (about \$7) and 34 sheqels (about \$7.7).

Results and discussion

The left-hand column of Figure 2 presents the observed maximization rate (selection of alternatives from class H) in the two problems as a function of the number of replicated alternatives. (The central column depicts the predictions under perfect chasing, and the right-hand column presents the predictions of the *post hoc* model to be discussed below.) The results reveal no significant linear relationship between the number of




Figure 2. Proportion of maximization, P(H), as a function of the number of replications (*m*) of each basic alternative in the two problems studied in Experiment 1. The notation N(*x*,*y*) implies a draw from a normal distribution with mean *x* and standard deviation *y*

alternatives and the maximization rate in Problem 1 [$t(19) = -0.23$, NS for the linear trend computed as the difference between $m = 1$ and $m = 25$]. The observed maximization rates are 0.71 (SD = 0.16), 0.71 (SD = 0.21), and 0.70 (SD = 0.26) with m levels of 1, 3, and 25, respectively. A different pattern was observed in Problem 2. The results in this problem show a significant negative linear relationship between the number of alternatives and the maximization rate [$t(19) = -5.98$, $p < 0.0001$]. The maximization rate in the binary choice condition (mean = 0.75, SD = 0.15) is higher than in both the 3-replications condition (mean = 0.55, SD = 0.2) and the 25-replications condition (mean = 0.44, SD = 0.23).

These results can be summarized as the combined effect of confusion and the tendency to chase returns on choice. We refer to this pattern as ‘confused chasing.’ In Problem 2 (where there is a higher variability in the lower mean), confusion and chasing work in the same direction. As the number of alternatives increases, the tendency to chase returns leads investors to prefer the riskier (and in this case inferior) prospect. Confusion leads investors toward random choice. Together, the two create a strong effect that shifts choices away from maximization. In contrast, in Problem 1, these two effects work in opposite directions and appear to cancel each other out.

EXPERIMENT 2: REPLICATED ALTERNATIVES AND RARE EVENTS

The second experiment was designed to examine the robustness of the joint effect of confusion and chasing to situations that involve rare events. In these situations, the implications of perfect chasing can be extremely sensitive to the number of replicated alternatives. The current experiment focuses on environments in which perfect chasing implies a strong tendency to select one class when $m = 1$ and a strong bias in the opposite direction when m is large. Under the assumption that the effect of confusion and chasing is additive, chasing is expected to have a larger effect than confusion in the current setting (relative to Experiment 1).

Experiment 2 focuses on the following basic problems:

Problem 3:

- H (EV = 3) A draw from the set {1,2,3,4,5}
- L (EV = 2.88) 32 with probability 0.09, 0 otherwise

Problem 4:

- H (EV = 3.2) 32 with probability 0.1, 0 otherwise
- L (EV = 3) A draw from the set {1,2,3,4,5}

Both problems study choice between a class of ‘long shot’ gambles and a class of safer prospects.

Previous studies on variants of Problem 4 with $m = 1$ (binary choice) show deviations from maximization in the direction of the perfect chasing predictions (see Barron & Erev, 2003; Yechiam & Bussemeyer, 2006; Yechiam et al., 2005b). These studies show that individuals behave as if they underweight rare outcomes. Accordingly, maximization rates are significantly below 50%.

Experiment 2 focused on the three m levels (number of replications) examined in Experiment 1. In Problem 3, perfect chasing implies maximization rates of 91%, 62%, and 10% for m levels of 1, 3, and 25, respectively. In this problem, both chasing and confusion imply lower maximization rates for higher levels of m . In Problem 4, however, perfect chasing implies maximization rates of 10%, 40%, and 92% for m levels of 1, 3, and 25, respectively. Here, both chasing and confusion imply higher maximization rates as m rises. Notice that in both problems the joint effect of confusion and chasing is expected to moderate the tendency of individuals to underweight rare positive outcomes (e.g., getting a relatively high payoff of 32). Yet the two hypotheses differ with respect to the magnitude of the predicted effect. Whereas confusion implies that the

maximization rate moves toward 50%, chasing (and an additive combination of confusion and chasing) implies a greater fall (toward 10%) in Problem 3, and a bigger rise (toward 92%) in Problem 4.

Design and procedure

The design and procedure were identical to Experiment 1, with the exception of the basic alternatives. Experiment 2 focused on Problems 3 and 4. Conversion rate in this experiment was 3 agorot (about 0.07 US cent) per point.

Participants

Twenty Technion students who did not participate in Experiment 1 served as paid participants in the experiment. Payment was made according to participants' obtained payoffs in the different conditions. Final payoffs ranged between 25 sheqels (\$4.7) and 30 sheqels (\$6.8).

Results and discussion

The results, presented in Figure 3 (using the same format as Figure 2), show a significant negative relationship between the number of alternatives and the maximization rate in Problem 3 [the difference between $m = 1$ and $m = 25$ is significant, $t(19) = -3.93, p < 0.001$]. The maximization rate in the binary choice condition (mean = 0.82, SD = 0.19) is similar to that in the 3-replications condition (mean = 0.81, SD = 0.32) and significantly higher than in the 25-replications condition (mean = 0.62, SD = 0.23). In Problem 4, however, the results reveal a significant positive relationship between the number of alternatives and the maximization rate [the difference between $m = 1$ and $m = 25$ is significant, $t(19) = 3.68, p < 0.001$]. The maximization rate in the binary choice condition (mean = 0.23, SD = 0.22) is much lower than in the 3-replications condition (mean = 0.39, SD = 0.32) and in the 25-replications condition (mean = 0.41, SD = 0.28).

The results suggest that the joint effect of chasing and confusion moderates the tendency of individuals to underweight rare positive outcomes. Consequently, the increase in the number of alternatives impairs maximization in Problem 3, but facilitates maximization in Problem 4. In addition, the results suggest that the joint effect of chasing and confusion cannot be captured with a simple additive model assuming that the observed behavior falls between the predictions of pure chasing and confusion (random choice). When $m = 3$ and 25 the observed maximization rates fall outside the interval between perfect chasing and random choice in both problems. For example, when $m = 25$ in Problem 4, perfect chasing implies a maximization rate of 92% while the observed rate is 40%. This pattern suggests that for large levels of m , the effect of confusion eliminates the effect of chasing.




Figure 3. Proportion of maximization, P(H), as a function of the number of replications (m) of each basic alternative in the two problems studied in Experiment 2. The notation (32, 0.1) implies a gamble that pays 32 with probability 0.1 and 0 otherwise. The notation {1,2,3,4,5} implies a gamble that pays each of the five outcomes with equal probability

EXPERIMENT 3—REGRET AND THE REVERSED CERTAINTY EFFECT

Experiment 3 was designed to compare two explanations for the large effect of confusion observed in the previous experiments. The first explanation suggests that the mere increase in the number of replicated alternatives is sufficient to cause strong confusion that can override the effect of chasing. A second interpretation is suggested by the results summarized in Erev and Barron (2005; see related ideas in Bell, 1982; Loomes & Sugden, 1982), and see related ideas in Bell (1982) and Loomes and Sugden (1982). Erev and Barron show that their results can be captured with the assumption that large forgone payoffs decrease payoff sensitivity. Their abstraction of this effect implies that random responses (i.e., the apparent ‘confusion’) are actually a product of experienced regret (information concerning higher forgone than obtained payoff).³ This assertion captures the current results because a rise in the number of replicated options (m) might increase the likelihood of regret. Specifically, the likelihood of regret, given a choice that maximizes EV, increases with m in Problems 1, 2, and 3 (but not in Problem 4 where a rise in m increased maximization). In Problem 3, for example, the exact values are 9%, 40%, and 92% with m levels of 1, 3, and 25, respectively.

In order to compare these explanations, Experiment 3 focuses on basic choice problems in which a rise in the number of replicated options has a distinct effect on the likelihood of regret. This experiment employs the same design as Experiments 1 and 2 to study the following basic problems:

Problem 5:

- H (EV = 34) A draw from the set {40, 41, 42, 43, 44} with probability 0.8
 A draw from the set {0, 1, 2, 3, 4} otherwise
- L (EV = 32) A draw from the set {30, 31, 32, 33, 34}

Problem 6:

- H (EV = 10) A draw from the set {40, 41, 42, 43, 44} with probability 0.2
 A draw from the set {0, 1, 2, 3, 4} otherwise
- L (EV = 9.5) A draw from the set {30, 31, 32, 33, 34} with probability 0.25
 A draw from the set {0, 1, 2, 3, 4} otherwise

Notice that the outcomes from Gamble H in Problem 5 and from the two basic gambles in Problem 6 are determined by two lotteries. The first lottery selects one of two sets, and the second selects one of the members of the selected set. This two-stage lottery design did not affect the experimental procedure. As in the previous studies the participants obtained the trial’s payoff immediately after each choice.

Problems 5 and 6 allow a comparison of the two explanations suggested above. Note that perfect chasing implies an increasing maximization rate with a rise in the number of alternatives in both problems (e.g., 80%, 99%, and 100% in Problem 5; 50%, 70%, and 99% in Problem 6 for m levels of 1, 3, and 25, respectively). Note further that the number of replicated alternatives in Problem 5 has a mild effect on regret. Ignoring the relatively small differences between the different members within each outcome set, the probability of regret (higher forgone payoff than the obtained payoff) given an H choice is less than 20%, independent of the number of replications.⁴ In Problem 6, however, the number of replicated alternatives has a strong effect on regret. In this problem, the probability of regret from choice H is 20% under binary choice, 58% under 3 replications, and 80% under 25 replications.

³Such information might make the decision maker regret her choices and decrease her sensitivity to the payoff structure. The focus on the maximal forgone payoff is also supported by some related studies of regret (Pyszczynski, Greenberg, & Laprelle, 1985; Tsiros, 1998).

⁴The assumption that small differences are ignored is presented here to simplify the presentation. The model, suggested below, abstracts regret in a slightly more complex way that avoids this simplification assumption.

Under the aforementioned analysis, the interpretation of regret predicts that in Problem 5 the chasing effect (increase in maximization with a rise in number of alternatives) will be stronger than the effect of confusion. However, in Problem 6 confusion will be stronger than chasing. The interpretation that raising the number of alternatives is sufficient to create confusion implies that in both problems confusion will be stronger than chasing.

Design and procedure

The design and procedure were identical to Experiments 1 and 2 with the exception of the basic alternatives. Experiment 3 focused on Problems 5 and 6. The conversion rate in this experiment was 0.5 agorot (about 0.012 US cent) per point.

Participants

Twenty Technion students who did not participate in Experiments 1 or 2 served as paid participants in the experiment. Payment was made according to participants' performance in the gambles. Final payoffs ranged between 19 sheqels (\$4.3) and 35 sheqels (\$7.95).

Results and discussion

The results, presented in Figure 4, reveal a positive linear relationship between the number of alternatives and the maximization rate in Problem 5. This trend (the difference between $m = 1$ and $m = 25$) is significant [$t(19) = 2.53, p < 0.02$]. The maximization rate in the binary choice condition (mean = 0.68, SD = 0.2) is lower than in the 3-replications condition (mean = 0.69, SD = 0.25) and in the 25-replications condition (mean = 0.80, SD = 0.26). This pattern is consistent with the prediction of the experienced regret explanation for the previous findings. It seems that the confusion effect—the bias toward random choice when the number of replicated options goes up—does not emerge when the increase in the number of options is not associated with an increase in experienced regret. In this case, the positive effect predicted by the chasing hypothesis is more robust.

The results in Problem 6 reveal only a weak relationship between the number of alternatives and the maximization rate [$t(19) = 1.94, NS$]. The maximization rate in the binary choice condition (mean = 0.42, SD = 0.27) is not significantly different from that in the 3-replications condition (mean = 0.51, SD = 0.1) or in the 25-replications condition (mean = 0.56, SD = 0.18).

Prob.	H	L
5	$(40, .8) + e$	$(30, 1) + e$
6	$(40, .2) + e$	$(30, .25) + e$




Figure 4. Proportion of maximization, P(H), as a function of the number of replications (m) of each basic alternative in the two problems studied in Experiment 3. The notation ' $(x, p) + e$ ' implies a payoff of $x + e$ with probability p , and a payoff of e otherwise. The value e is randomly selected (independently for each option) from the set $\{0,1,2,3,4\}$

Notice that Problems 5 and 6 are variants of the problems used by Kahneman and Tversky (1979), following Allais (1953), to demonstrate the certainty effect. Problem 6 was created from Problem 5 by dividing the probability of the desirable set by 4. As noted by Allais, this linear transformation is not expected to affect preference under expected utility theory. Yet when people are faced with these problems once, and can rely on a description of the possible payoff distributions, they prefer L in variants of Problem 5, and H in variants of Problem 6. That is, the safer option (L) is more attractive in when it provides the desirable payoff with certainty (in variants of Problem 5). Kahneman and Tversky (1979) referred to this pattern as the certainty effect. Barron and Erev (2003) (see also Hertwig, Barron, Weber, & Erev, 2004) show that the opposite pattern emerges in decisions from experience. They observed less choices of the safer option in a variant of Problem 5 (37% L choices) than in a variant of Problem 6 (49% L choices). The current results demonstrate the robustness of the reversed certainty effect. The safer option was less attractive in Problem 5 (when it provided the desirable payoff with certainty) than in Problem 6 in all three m levels: [$t(19) = 3.28$, $p < 0.005$; $t(19) = 3.36$, $p < 0.005$; and $t(19) = 5.32$, $p < 0.001$, with $m = 1, 3$, and 25 , respectively].

A QUANTITATIVE SUMMARY: THE 'CHASING AND REGRET' MODEL

The current results can be summarized with the assertion that the number of replicated alternatives affects decisions from experience because it influences two factors important to the learning process: the properties (payoff distribution) of the alternatives with the best recent payoffs, and experienced regret. This verbal summary of the results is naturally quantified with some of the models studied in Erev and Barron (2005). The simplest model in this set is a generalization of the basic Fictitious Play model (see Brown, 1951 and review in Camerer, 2003) that allows for the possibility that regret reduces payoff sensitivity. This model, referred to here as the 'chasing and regret' model, can be summarized with the following assumptions:

A1: Weighted adjustment.

The adjusted propensity to select alternative j at trial $t + 1$ is:

$$q_{j,t+1} = (1 - w)q_{jt} + w \cdot x_{jt} \quad (1)$$

where x_{jt} is the payoff of j in trial t , and w is a parameter that determines the weight of this payoff. The initial value is $q_{jt} = 0$.

A2: Stochastic response with sensitivity to experienced regret.

The probability that the decision maker will select strategy j in trial t is:

$$p_{jt} = \frac{e^{q_{jt} \cdot \lambda / s_t}}{\sum_{k=1}^n e^{q_{kt} \cdot \lambda / s_t}} \quad (2)$$

where λ is a parameter that determines the relative importance of the different propensities, and S_t is the experienced regret level of the decision maker, determined by:

$$S_{t+1} = (1 - w)S_t + w|\max_t - x_{jt}| \quad (3)$$

where \max_t is the maximal payoff obtained in trial t over the $2m$ alternatives. The initial regret level is set to equal $S_1 = \lambda$.

Descriptive value

To test if the chasing and regret model can capture the current results, we ran computer simulations in which virtual agents that behave in accordance with the model's predictions participate in the 18 experimental

conditions. The simulations were run with different sets of parameters, in an effort to find those that minimize the mean squared deviation (MSD) between the model's prediction and the observed maximization rates (presented in the left-hand columns of Figures 2–4). The mean squared deviation between the data and the predictions was minimized when the two free parameters were set to $w = 0.52$, and $\lambda = 2.75$. The model's predictions with these parameters are presented in the right-hand columns of Figures 2–4. The figures show that the model captures the main trend of the results. Over the 18 conditions, the correlation between the observed and the predicted proportions is 0.90. The MSD score is 0.0013.

Learning patterns

Figure 5 presents the observed and predicted maximization rates in five blocks of 10 trials in each of the 18 experimental conditions. The experimental results reveal relatively flat learning curves. Indeed, in only 7 of the 18 cases did the observed rate increase from the first to the last block. The predictions of the model show a similar pattern.

Sequential dependencies

Figure 6 presents two measures of the dependencies between sequential trials. The first measure is the 'chasing rate': the proportion of times in which the option with the highest return in trial t was selected in trial $t + 1$. The second measure, referred to as 'repetitions,' is the proportion of times in which the alternative selected in trial t was selected again in trial $t + 1$.

These results show that both proportions exceed random choice in all conditions. In addition, the results reveal an interesting nonlinear relationship between the number of replicated alternatives and the repetition rate. It seems that the increase from 3 to 25 replicated alternatives had very little effect on the repetition rate.

The right-hand column of Figure 6 shows that the chasing and regret model captures the observed chasing rate, but fails to capture the nonlinear repetition rate. This failure can be addressed by modifying the choice rule to allow for the possibility that with probability q , the decision maker simply repeats his or her last choice (see Erev & Haruvy, 2005). Another possibility is that the decision makers handle a presentation with large set of alternatives by limiting their attention to a subset of these alternatives. Such a tendency facilitates repetitions and can be captured with the addition of an attention parameter to the current model.

More complex models

Erev and Barron (2005) show that the chasing and regret model is outperformed by a four-parameter model referred to as RELACS (reinforcement learning among cognitive strategies). RELACS assumes that the decision maker learns among three strategies that can be referred to as 'stochastic best reply,' 'deterministic best reply,' and 'case-based reasoning.' The chasing and regret model is a simplification of RELACS in which only one of the three strategies, stochastic best reply, is used. A comparison of RELACS with the chasing and regret model on the current data reveals that in the current setting the added complexity assumed by RELACS is not necessary. The two-parameter simplification is as accurate as the four-parameter model⁵.

A look at the difference between the relative value of the simple chasing and regret model here and in Erev and Barron (2005) suggests that the current findings arise from the effect of the available feedback. The chasing and regret model provides a very good approximation of learning in situations with complete feedback (here and in Erev & Barron, 2005). This model is much less useful in accounting for learning in

⁵RELACS was designed to address binary choice tasks. It has to be generalized to address the current tasks. Under one generalization that captures the current results the decision maker updates only two alternatives after each choice: the selected option, and the option that led to the best payoff in the previous trial. The MSD of this generalization is practically identical to the MSD of the simpler chasing and regret model.

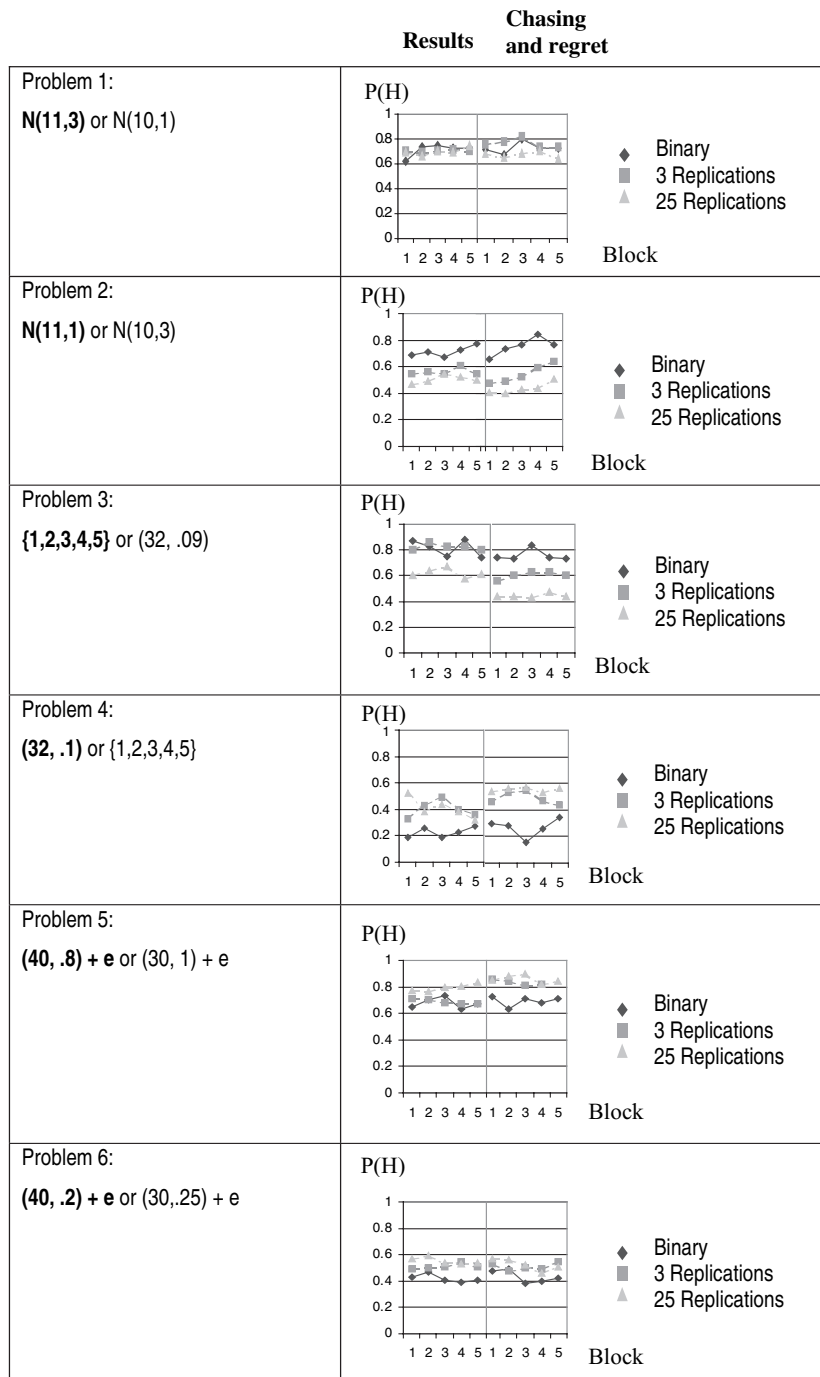


Figure 5. Observed and predicted learning curves: $P(H)$ in blocks of 10 trials, by problem and number of replications (m). The left-hand columns present the experimental results. The right-hand columns present the predictions of the chasing and regret model

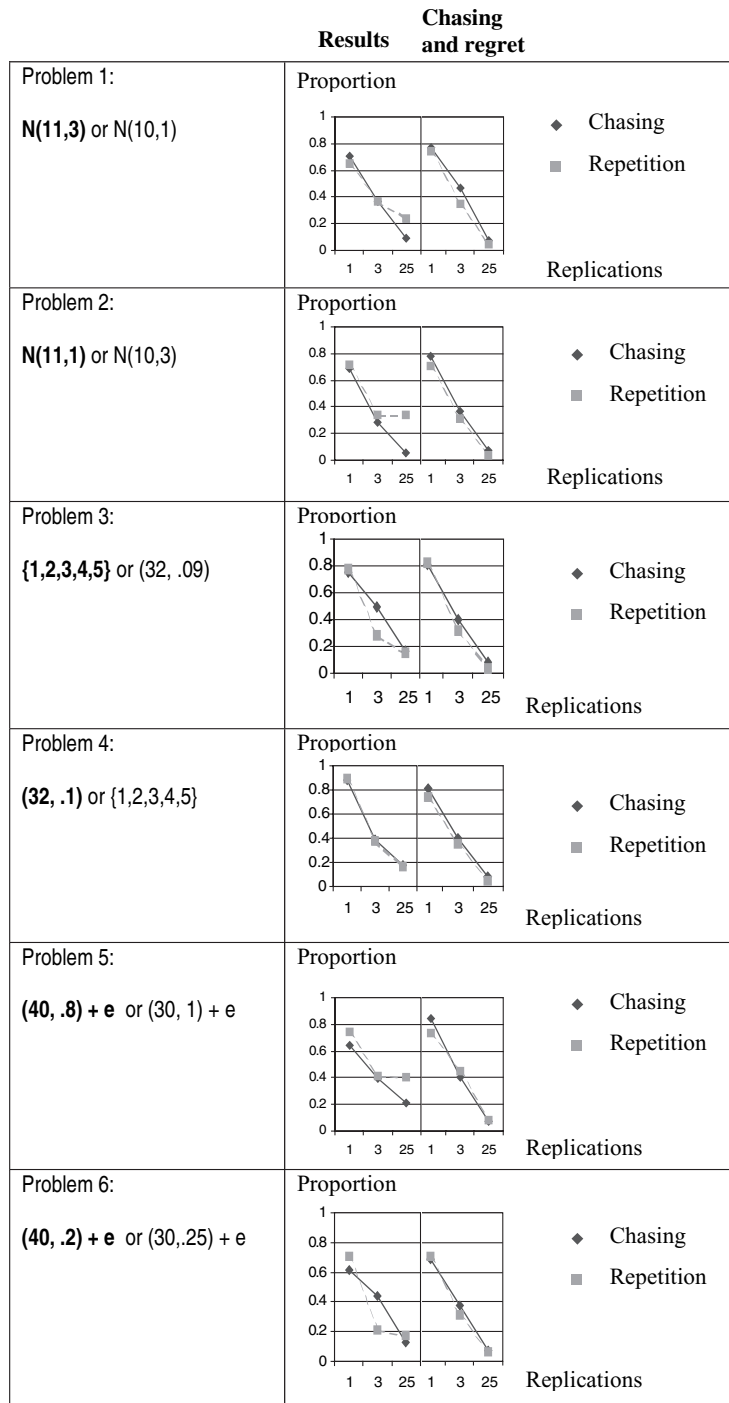


Figure 6. Chasing rate (selecting the alternative with the best payoff in the most recent trial) and repetition rate (repeating most recent choice) by problems and number of replications (m). The left-hand columns present the experimental results. The right-hand columns present the predictions of the chasing and regret model

environments where the available feedback after each choice is restricted to the obtained payoff. The more complex process assumed by RELACS is necessary in order to capture learning in situations with restricted feedback.

GENERAL DISCUSSION

The current paper explores decision making between alternatives that belong to two distinct classes. Previous studies of these environments have focused on the effect of asymmetry in the number of alternatives in each class (e.g., Benartzi & Thaler, 2001). That research shows a bias toward the class with more members. The three experiments presented here examine the symmetrical case, exploring the effect of a constant increase in the number of replicated alternatives in the different classes. The results show that when the prior information available to the decision maker is limited, the increase in the number of replications has interesting effects.

Experiment 1 represents an environment in which a rise in the number of replicated alternatives increases the attractiveness of the class associated with a lower expected return and higher risk. These results can be summarized with the assertion that replicated alternatives create confusion and change the implications of the tendency to chase recent returns. Experiment 2 presents a situation in which a rise in the number of replicated alternatives increases the implied weighting of rare events. Experiment 3 suggests that the relative strengths of chasing and confusion effects are determined by experienced regret. It demonstrates that this understanding can be applied to enhance maximization in variants of the Allais paradox.

The main experimental results can be summarized with a simple model, which assumes that payoff sensitivity decreases with experienced regret. This 'chasing and regret' model relies on the observation that an increase in the number of replicated alternatives has two objective effects. First, it tends to increase the experienced regret (the difference between the obtained payoff and the ex-post best payoff). Second, it tends to increase the probability that one of the members of the riskier class will yield the best payoff. The model captures the main results with the assumption of a simple cognitive process whose output is sensitive to these two objective effects.

The role of information concerning the underlying categories

It is important to recall that the current investigation focuses on situations in which the decision makers are not informed that the different alternatives can be classified into two classes. Examinations of situations in which the participants are informed that the alternatives belong to two classes reveal that this information can erase the effect of the number of alternatives. For example, Ert (in preparation) ran a variant of Experiment 2 in which the difference between the two classes was highlighted with distinct colors. The participants were informed that all the keys with the same color are associated with the same payoff variability level. The experiment was identical to Experiment 2 in all other respects. The observed maximization rates were not affected by the number of alternatives. The observed means in Problem 3 were 0.71, 0.70, and 0.71 for m values 1, 3, and 25; the means in Problem 4 were 0.34, 0.36, and 0.32 for m values 1, 3, and 25, respectively.

These results suggest that the confused chasing pattern can be reduced by reliable information concerning the distinct classes. It seems that in these cases the assumption that decision makers choose among the basic classes can be a reasonable simplification of the choice task. We return to this suggestion below in the discussion of the practical implications of the results.

Relationship to the detrimental effect of choice

In many natural multi alternative choice problems, the decision makers can decide to reject all the alternatives. For example, supermarket shoppers can select one of many jams, or not to buy jam at all.

Examination of this setting (e.g., Boatwright & Nunes, 2001; Iyengar & Lepper, 2000) led to the interesting observation that an increase in the number of choice alternatives (e.g., distinct jams) can reduce the probability that one of the alternatives will be selected.

The common explanation for this 'detrimental effect of choice' assumes that adding choices increases choosers' confusion (Huffman & Kahn, 1998; Iyengar & Lepper, 2000) and leads to weaker preferences (Chernev, 2003; Dhar, 1997; Gourville & Soman, 2005; Iyengar & Lepper, 2000). The current research highlights another possible contributor to this effect. The chasing and regret pattern implies that under certain settings an increase in the number of alternatives reduces the probability that the best alternative will be selected. Reaction to this effect can increase the tendency to avoid the choice among the different alternatives when the number of alternatives increases.

Practical implications

We believe that the most important practical implications of the current analysis involve the observation that an increase in the number of alternatives increases risk seeking in pure decisions from experience. This effect is particularly large when the risky alternatives are counterproductive (associated with low EV), and it emerges even when all the alternatives are replicas of two basic gambles.

One set of implications of this observation involves the design of decision environments that enhance safety. For example, our findings imply that rules and constraints that reduce the number of acceptable behaviors can be effective in bolstering safety even if they do not change the ratio of safe and risky behaviors. The mere reduction in the number of acceptable behaviors (both risky and safe) is predicted to result in more people choosing safe over risky alternatives. A second intervention that can reduce risk seeking in the current context is reliable categorization of the different behaviors. As noted above, clear information concerning the distinct categories can eliminate the effect of the number of alternatives.

The same observation can also be used to facilitate the design of decision environments that enhance risky behavior. It seems that the value of this suggestion is well known to casino designers. They apply a similar idea by investing money and space in order to increase the number of replicated options like slot machines.

A third set of implications involves investment decisions. The current results provide a possible explanation for the observation that the typical investor appears to avoid risky investments while selecting between bonds and stocks (see Mehra & Prescott, 1985; Thaler, Tversky, Kahneman, & Schwartz, 1997), but tend to exhibit risk seeking (e.g., by holding poorly-diversified portfolios) while comparing different stocks portfolios (e.g., Barber & Odean, 2000; Blume & Friend, 1975; Kelly, 1995). Under the current explanation this pattern reflects the fact that in the first decision the distinction between the two categories (stocks and bonds) is explicit. Thus, the large number of alternatives has little effect. The distinction between risky and safe prospects is much less obvious during the selection among stocks portfolios (and/or mutual funds). In this case, the implied risk is a function of the correlation between the different investments. Thus, this behavior is more sensitive to the confused chasing effect discussed here.

A final set of implications involves the robustness of the tendency to underweight rare events in decisions from experience (see Barron & Erev, 2003; Erev & Barron, 2005; Hertwig et al., 2004). The current results suggest that when the rare events are unattractive (like accidents, terrorist attacks, and the outcome of zero in Problem 5), this tendency is enhanced by an increase in the number of alternatives. This suggestion is consistent with the observation that the terrorist attacks have larger effect on the (binary) decision whether to visit an area under attack, than on the selection among the different activities in that area (see Yechiam et al., 2005a). An increase in the number of alternatives appears to increase the effect of rare attractive outcomes (low probability prizes and the 32 outcome in Problem 4). However, this increase is relatively small (see Problem 4), and is eliminated by the provision of information concerning the different categories.

Summary

The current analysis suggests that an increase in the number of replicated alternatives has two objective effects that trigger interesting behavioral phenomena. The first objective effect involves an increase in the probability that one of the riskier alternatives will yield the best payoff. The results suggest that human behavior is sensitive to this effect because people tend to select one of the alternatives with highest recent payoffs. The second objective effect involves an increase in experienced regret (the difference between the obtained payoff and the *ex-post* best payoff). The results suggest that an increase in experienced regret reduces payoff sensitivity. The joint effect of these tendencies can be captured with a simple ‘chasing and regret’ model.

APPENDIX—INSTRUCTION

In this experiment you are operating a computerized money machine. Upon pressing a button you will win or lose some points (the button’s ‘payoff’ in that trial). Your goal is to collect as many points as possible. The buttons may be different from each other. Your final payment will be determined by the number of points you earn (1 point = 1 agora). For your information, the machine is likely to differ between participants. Good luck.⁶

ACKNOWLEDGEMENTS

The authors thank Frank Yates and three anonymous reviewers for their valuable suggestions, and Meira Ben-Gad for editorial assistance. This research was supported by a grant from USA-Israel Binational Science Foundation.

REFERENCES

- Allais, M. (1953). Le comportement de l’homme rationnel devant le risque, critique des postulats et axiomes de l’école Américaine. *Econometrica*, *21*, 503–546.
- Arrow, K. J. (1959). Rational choice functions and orderings. *Economica*, *26*, 121–127.
- Barber, B. M., & Odean, T. (2000). Trading is hazardous to your wealth: The common stock investment performance of individual investors. *The Journal of Finance*, *55*, 773–806.
- Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description based decisions. *Journal of Behavioral Decision Making*, *16*, 215–233.
- Bell, D. E. (1982). Regret in decision-making under uncertainty. *Operations Research*, *30*, 961–981.
- Benartzi, S., & Thaler, R. (2001). Naïve diversification strategies in defined contribution saving plans. *American Economics Review*, *91*, 79–98.
- Blume, M., & Friend, I. (1975). The asset structure of individual portfolios and some implications for utility theory. *Journal of Finance*, *30*, 585–603.
- Boatwright, P., & Nunes, J. C. (2001). Reducing assortment: An attribute based approach. *Journal of Marketing*, *65*, 50–63.
- Brown, G. W. (1951). Iterative solutions of games by fictitious play. In T. C. Koopmans (Ed.), *Activity Analysis of Production and Allocation*. New York: Wiley.
- Budescu, D. V., Erev, I., Wallsten, T. S., & Yates, J. F. (1997). Introduction to a special issue on stochastic and cognitive models of confidence. *Journal of Behavioral Decision Making*, *10*, 153–155.

⁶We present the instructions of the first experiment. The instructions for Experiment 2, and 3 were the same with the exception of the conversion rates that differed between experiments.

- Camerer, C. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton, New York: Princeton University Press/Russell Sage Foundation.
- Chernev, A. (2003). When more is less and less is more: The role of ideal point availability and assortment in consumer choice. *Journal of Consumer Research*, *30*, 170–183.
- Dhar, R. (1997). Consumer preference for a no-choice option. *Journal of Consumer Research*, *25*, 369–384.
- Elrod, T., Louviere, J. J., & Davey, K. S. (1992). An empirical comparison of rating based and choice-based conjoint models. *Journal of Marketing Research*, *29*, 368–377.
- Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review*, *112*, 912–931.
- Erev, I., & Haruvy, E. (2005). Generality, repetition, and the role of descriptive learning models. *Journal of Mathematical Psychology*, *49*(5), 357–371.
- Ert, E. (in preparation). Learning in a multi alternative environment. PhD Dissertation, Technion—Israel Institute of Technology.
- Ferster, C. B., & Skinner, B. F. (1957). *Schedules of reinforcement*. East Norwalk, CT, USA.
- Goldstein, W. M., & Mitzel, H. C. (1992). The relative importance of relative importance: Inferring other people's preferences from relative importance ratings and previous decisions. *Organizational Behavior and Human Decision Processes*, *51*, 382–415.
- Gourville, J. T., & Soman, D. (2005). Overchoice and assortment type: When and why variety backfires. *Marketing Science*, *24*, 382–395.
- Grosskopf, B., Erev, I., & Yechiam, E. (2006). Forgone with the wind: Indirect payoff information and its implications for choice. *International Journal of Game Theory*, *34*, 285–302.
- Hellriegel, D., Slocum, J. W., & Woodman, R. W. (1998). *Organizational behavior* (8th ed.). South Western: Cincinnati, Ohio, USA.
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choices. *Psychological Science*, *15*, 534–39.
- Hong, H., Lim, T., & Stein, J. C. (2000). Bad news travel slowly: Size, analyst coverage and the profitability of momentum Strategies. *The Journal of Finance*, *55*, 265–295.
- Hong, H., & Stein, J. C. (1999). A unified theory of under-reaction, momentum trading and overreaction in asset markets. *The Journal of Finance*, *54*, 2143–2184.
- Huberman, G., & Jiang, W. (2006). Offering vs. choice in 401(k) plans: Equity exposure and number of funds. *Journal of Finance*, *61*, 763–801.
- Huffman, C., & Kahn, B. E. (1998). Variety for sale: Mass customization or mass confusion? *Journal of Retailing*, *74*, 491–513.
- Iyengar, S. S., & Lepper, M. R. (2000). When choice is demotivating: Can one desire too much of a good thing? *Journal of Personality and Social Psychology*, *79*, 995–1006.
- Jacoby, J., Speller, D. E., & Kohn, C. A. (1974). Brand choice behavior as a function of information load. *Journal of Marketing Research*, *11*, 63–69.
- Jegadeesh, N., & Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance*, *48*, 65–91.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263–291.
- Keller, K. L., & Staelin, R. (1987). Effects of quality and quantity of information on decision effectiveness. *Journal of Consumer Research*, *15*, 504–508.
- Kelly, M. (1995). All their eggs in one basket: Portfolio diversification of US households. *Journal of Economic Behavior and Organization*, *27*, 87–96.
- Loomes, G., & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal*, *92*, 805–824.
- Luce, R. D. (1959). *Individual choice behavior: A theoretical analysis*. New York: Wiley.
- Mehra, R., & Prescott, E. C. (1985). The equity premium puzzle. *Journal of Monetary Economics*, *XV*, 145–161.
- Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, *63*, 81–97.
- Nash, J. F. (1950). The bargaining problem. *Econometrica*, *18*, 155–162.
- Pyszczynski, T., Greenberg, J., & Laprelle, J. (1985). Social comparison after success and failure: Biased search for information consistent with a self-serving conclusion. *Journal of Experimental Social Psychology*, *21*, 195–211.
- Rouwenhorst, K. G. (1998). International momentum strategies. *Journal of Finance*, *53*, 267–284.
- Shavit, T., Ben Zion, U., Erev, I., & Haruvy, E. (2005). Diversification and the role of recent returns. Working paper: Technion.

- Thaler, R., Tversky, A., Kahneman, D., & Schwartz, A. (1997). The effect of myopia and loss aversion on risk taking: An experimental test. *The Quarterly Journal of Economics*, *112*, 647–661.
- Tsiros, M. (1998). Effect of regret on post-choice valuation: The case of more than two alternatives. *Organizational Behavior and Human Decision Processes*, *76*, 48–69.
- Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, *79*, 281–299.
- Tversky, A., & Simonson, I. (1993). Context-dependent preferences. *Management Science*, *39*, 1179–1189.
- Yechiam, E., & Busemeyer, J. (2006). The effect of forgone payoffs on underweighting small probability events. *Journal of Behavioral Decision Making*, *19*, 1–16.
- Yechiam, E., Barron, G., & Erev, I. (2005a). Description, experience and the effect of rare terrorist attacks. *Journal of Conflict Resolutions*, *49*, 430–439.
- Yechiam, E., Stout, J. C., Busemeyer, J. R., Rock, S. L., & Finn, P. R., (2005b). Individual differences in the response to forgone payoffs: An examination of high functioning drug abusers. *Journal of Behavioral Decision Making*, *18*, 97–110.

Authors' biographies:

Eyal Ert is a Ph.D. candidate in Behavioral Sciences at the Technion. His current research interests focus on individual learning and decision making, and their implications to marketing strategies and consumer behavior. Other interests include the implications of learning to behavioral game theory.

Ido Erev is the ATS' Women's Division Professor of Industrial Engineering and Management at the Technion. His current research focuses on modelling and deriving the implications of human adaptation to economic incentives. Examples of implications include the design of rule enforcement and bonus systems, training methods, and marketing strategies.

Authors' Addresses:

Eyal Ert and **Ido Erev** The Max Wertheimer Minerva Center for Cognitive Studies, Faculty of Industrial Engineering and Management, Technion, Haifa 32000, Israel.