The effect of foregone payoffs on underweighting small probability events

Eldad Yechiam

Technion - Israel Institute of Technology

Jerome R. Busemeyer

Indiana University, Bloomington IN

RUNNING HEAD: Foregone payoffs and rare events

Mail should be addressed to the first author:

Eldad Yechiam, Max Wertheimer Minerva Center for Cognitive Studies, Faculty of Industrial Engineering and Management, Technion - Israel Institute of Technology, Haifa 32000, Israel. Phone: (972) 4-829-4420 , Fax: (972) 4-829-5688, Email: eyehiam@indiana.edu.

The authors would like to thank Katie Davis and Tiffany Glist for their assistance with collecting data. We gratefully acknowledge support from the Center for the Study of Institutions, Population, and Environmental Change at Indiana University through National Science Foundation grants BCS-0215738 and SES0083511.

Abstract:

Foregone payoffs add information about the outcomes for alternatives that are not chosen. The present paper examines the effect of foregone payoffs on underweighting rare but possible events in repeated choice tasks. Previous studies have not demonstrated any long-lasting effects of foregone payoffs (following repeated presentation of a task) when foregone payoffs do not add much information. The present paper highlights the conditions and the contributing factors for the occurrence of such long-lasting effects. An experimental study compares the effect of foregone payoffs under different degrees of rarity of the negative payoff. It is demonstrated that foregone payoffs increase the selection from risky alternatives with extremely rare and highly negative outcomes, and that this effect does not diminish with repeated presentation of the task. These findings can be summarized using a surprisingly simple reinforcement-learning model. The findings are discussed in the context of the potential long-term effect of social learning.

Keywords: Decision making, experience, learning, foregone payoffs, social learning

In experience-based decisions, decision makers learn the relevant payoff distributions from trial by trial feedback. Recent experimental findings indicate that *experience* leads to underweighting of rare but possible events relative to decisions based on descriptions of the gambles (Barron & Erev, 2003; Hertwig, Barron, Weber & Erev, 2004; Weber, Shafir & Blais, 2004; Yechiam, Barron & Erev, 2005). For example, Barron and Erev (2004) showed that when repeatedly selecting from two options: One producing 3 Israeli Agora for sure and another option producing 32 Agora 10% of the time (and zero otherwise), decision makers tend to make more choices of the former gamble compared to when they are presented with the gambles' description, behaving as if they are giving less weight to the small probability event (of gaining 32 Agora). The goal of the present paper is to examine the effect of a second variable on the weighting of rare events in experience-based decisions. We evaluate the effect of presenting foregone payoffs, that is information about the outcomes for alternatives that are not chosen. This type of payoff, it is argued, enhances the underweighting of small probability events.

Some studies report that when foregone payoffs do not add much information about the potential outcome (see below for an elaboration), the effect of foregone payoffs is transitory and disappears with experience, that is, with repeated presentations of the task, (Grosskpf, Erev, & Yechiam, 2005; see also studies of games: Charness & Grosskopf, 2004; Grosskopf, 2003), and a study by Haruvy and Erev (2002) has demonstrated no significant effects of foregone payoffs (analyzed in Grosskopf et al., 2005). The present study highlights an important factor that leads to a long lasting effect of foregone payoffs on risk taking in experience-based choice behavior: The rareness of the negative outcomes.

Note that a relatively trivial effect of foregone payoffs in experience-based choice problems is due to the fact that foregone payoffs can increase the amount of information (Haruvy and Erev, 2002; Grosskopf et al., 2005). This can lead to long lasting effects when the choice problem has many options that require exploration in order to find out the optimal outcome. Presenting foregone payoffs implies giving full information about the outcomes from all of the alternatives, and reducing the need for exploration[1]. We focus on less trivial effects of foregone payoffs. Namely, we argue that even though decision makers can easily gather information about the potential outcomes (e.g., in binary choice tasks), there could be situations where foregone payoffs would have enduring and significant impacts on risk taking.

Effects of foregone payoffs in experience based choice:

The effect of anticipating foregone payoff information has been extensively assessed in the judgment and decision making literature (e.g., Josephs, Larrick, Steele & Nisbett, 1992; Zeelenberg, 1999). These studies have been instrumental in establishing that people are sensitive to the availability of foregone payoffs. However, these previous studies focused on knowing that one would get foregone payoffs (i.e., expectations about feedback type), rather than on the effect of actually getting foregone payoffs information on risk-taking in subsequent decisions.

[1] For example, consider a person who chooses repeatedly in a matrix of 20 × 20 cells which has randomly distributed outcomes consecutively running from 1 to 400 (i.e., each button has an attached outcome which is either 1, 2, 3 etc.). These outcomes are not known initially but are revealed after each choice. Obtaining foregone payoffs leads to information about the past outcomes of all of the cells *in a single trial*. When the payoff environment has strong local optima, without foregone payoffs intuitive exploration may to converge in local optima (see Yechiam, Erev & Gopher, 2001; Erev & Barron, 2005). Thus, foregone payoffs can have long lasting effects as they clearly point out the location of the global optimum (400) in the first trial (see Grosskopf et al., 2005)[1].

The effects of actually getting foregone payoff information on risk taking were previously investigated by Haruvy and Erev (2002). They evaluated the effect of foregone payoffs in two repeated choice problems of 200 trials. Decision makers were asked to choose between two buttons that appeared on the screen, a safe alternative (no noise) and a more risky alternative. The payoffs for pressing each button were as follows (1 Israeli Agora equaled about 0.25 cents at the time):

	Option R (risky)	Option S (safe)
Problem 1.	+1 Agora with p=.5; +21 otherwise	+10 Agora with certainty
Problem 2.	-1 Agora with p=.5; -21 otherwise	-10 Agora with certainty

The two problems were presented with or without foregone payoffs. Foregone payoff was presented on the unselected button but was not added to the tally below the button indicating the earning on the current trial, and to the accumulated payoff tally (see figure 1; which demonstrates the current experiment that uses the same design).

Note that option R (the more risky high-variability option) has higher expected value (10.5 compared to 10) in the gain domain and lower expected value (-10.5 compared to –10) in the loss domain. The results showed a small (about 5% difference) and not statistically-significant effect of foregone payoffs (see Grosskopf et al., 2005). The direction of the effect was in both conditions towards the option with the higher expected value (option R in the gain domain, option S in the loss domain). Clearly though, decision makers who were presented with foregone payoffs did not move towards the risky or the safe option.

Grosskopf et al. (2005) suggested that the effect of foregone payoffs on risk taking would increase with the number of alternatives because in a multiple-alternative problem decision makers are continually presented with the high spectrum of the outcome of risky alternatives. To examine their prediction, Grosskopf et. al., (2005) studied repeated choice behavior among 100 alternatives arranged in a matrix, where decision makers did not initially know the underlying probabilities and payoffs. The payoffs were actually drawn from one of two normal distributions: Fifty S (safe) buttons were drawn from N~(11,1), and 50 R (risky) buttons were drawn from N~(10,3). The alternatives appeared as one hundred blank buttons; there was no way of telling which button was related to the two distributions. The arrangement of the S and R buttons in the matrix was randomized for each player. Half of the participants performed the task without foregone payoffs, and the other half were presented with foregone payoffs for all of the (99) options that they had not chosen in addition to their selection. The results showed that in the foregone payoff condition the proportion of choices from R buttons increased (from .41 to .53) in the first 100 trials. This effect has been termed by Grosskopf et al. (2005) “the big eyes” (following Erev & Rapaport, 1998) because it was attributed to the fact that under a foregone payoffs condition participants may imagine better the potential gains from risky alternatives. For example, in a typical trial, some R alternatives would present high outcomes (e.g., 15 or 16) that may draw decision makers away from S alternatives that typically produce 9, 10, or 11.

The effect of foregone payoffs observed in Grosskopf et al., (2005) was limited however to the initial half of the task. In the second half the choice proportion of R was similar in both conditions (average .50 in both). Thus, on average, the effect of foregone payoffs diminished with experience. Grosskopf et al. (2005) suggested that this was because decision makers learned to ignore the foregone payoffs and revert back to choosing safer options, and that the effect is transitory[2].

We note that both in Haruvy and Erev (2002) and in Grosskopf et. al., (2005) the risky alternative contained a negative outcome (compared to the typical outcomes of the safe alternative) that appeared about 50% of the time or more. This is also true to studies of foregone payoffs in games which focused on situations where the risky alternative does not produce better results most of the time (see e.g., Charness & Grosskopf, 2004). We focus on risky alternatives that produce high payoffs most of the time but also *relatively rare and highly negative penalties*. We argue that foregone payoffs increase the attractiveness of such risky alternatives, and that decision makers do not learn to ignore foregone payoffs. Such an effect is interesting as it implies that more information (foregone payoffs) leads people to make riskier choices, and that people essentially cannot learn to use the additional information to make safer choices even after many repetitions of the task.

The effect of foregone payoffs and rare negative events

Consider for example the following two simplified choice problems. In each problem decision makers have to repeatedly select between two gambles:

[2] However, it could be that the choice of R stabilized at .50 (which could reflect a random allocation of selections).

Problem 3.	S (Safe)	Lose 8 cents with a probability of 0.005 (1 in 200)
		Lose 2 cents otherwise
	R (Risky)	Lose 300 cents with a probability of 0.005 (1 in 200)
		Lose 1 cent otherwise

Problem 4.	S (Safe)	Lose 8 cents with a probability of 0.005 (1 in 200)
		Lose 2 cents otherwise
	R (Risky)	Lose 30 cents with a probability of 0.05 (1 in 20)
		Lose 1 cent otherwise

In both problems the expected value of option R is the same (-2.495), and in both problems it is approximately 25% lower than in option S (-2.03). Moreover, in both problems, option R has higher variability (SD=21 in problems 3; 6 in Problem 4) and is thus considered to be more risky than option S (SD=0.4). This is the case because option R includes a rare possibility (1/20 or 1/200) of a relatively large loss. Yet in Problem 3 there is a less frequent chance (1/200 compared to 1/20) of losing a larger amount (300 compared to 30 cents). Thus, the two problems are different in the degree of rareness of the negative outcome while the size of the outcome is linearly compensated by the same degree (i.e., the expected value is held constant). For the ease of reading, these problems are re-dubbed as 1/200 (problem 3) and 1/20 (problem 4).

We predicted that the effects of foregone payoffs would be stronger in problems such as 1/20 and 1/200 than in problems like 1-2 for the following reason. In these problems when decision makers choose the safe alternative they usually get an outcome of –2 (or very rarely –8), and most of the times (199 out of every 200 trial in Problem 1/200 and 19 out of every 20 trials in problem 1/20) the foregone outcome in the risky alternative is only -1. It has been previously demonstrated (by Estes, 1976a; 1976b; Erev, Haruvy & Yechiam, 2003) that people are more sensitive to what happens most of the time than to the average outcome. This implies that decision makers in both problems would be attracted to the foregone outcomes of the risky alternative because they are better most of the time compared to the outcomes from the safe alternative that are better on the average. Moreover, the effect in both conditions is not likely to diminish over time, as the tendency to select the alternative that produces better outcomes most of the time is not assumed to change (see Erev et al., 2003). Finally, even though the average foregone outcome in the risky alternative is the same in problems 1/20 and 1/200, the attraction to the risky alternative's foregone payoffs would be larger in problem 1/200 because in this problem *on the vast majority of trials,* foregone payoffs indicate better outcomes (-1) for the risky alternative.

The notion that "most of the time is better than the average" is embedded in a surprisingly simple formal reinforcement-learning model, which combines a Delta learning rule (see Busemeyer & Myung, 1992; Gluck & Bower, 1988; Sutton & Barto, 1998; Sarin and Vahid, 1999) with partial weighting of foregone payoffs (see Camerer & Ho, 1999; Yechiam & Busemeyer, 2005a). This occurs in the model because of exponential discounting of foregone outcomes. In other words, people react based on recent events and they quickly discount past outcomes. Thus, for example, in the 1/200 problem, the recent foregone outcomes from the risky alternatives are almost always better than the outcomes from the safe alternative, leading to the attraction to the risky alternative in this problem.

The model is described fully in Appendix 1. This model was run using a simulation based on parameters that were previously estimated in a different choice task (Yechiam & Busemeyer, 2005a; see Appendix for details). One thousand simulations were generated to produce a distribution of 400 repeated choice sequences for problem 1/200 and 1/20. These distributions (which form the predictions for the results in experiment 1) appear in Figure 2. Based on the estimated parameters, the model predicts that (a) in terms of overall proportion of risky choices predicted, in both problems foregone payoffs increase the proportion of R choices, and the effect is much stronger in Problem 1/200 than in Problem 1/20; (b) in terms of changes in this pattern over time, the effect of foregone payoffs on risk taking increases following repeated selections, and this increase is larger in Problem 1/200. The more complex Reinforcement Learning Among Cognitive Strategies (RELACS) model of Erev and Barron (in press) reproduces these results as well.

<u>Experiment: Long lasting effects of foregone payoffs</u>

To examine the differential effects of foregone payoffs under different payoffs conditions, we presented the two problems above (1/200 and 1/20) experimentally in a repeated-choice task. Participants were asked to repeatedly select between two buttons for an unknown number of repetitions. Half of the participants were presented with no foregone payoffs (Foregone-0 condition). In this condition players only saw the payoffs from the selected buttons. The other half were presented with foregone payoffs once every two trials (Foregone-1 condition). Namely, in these trials they received full information about the payoffs from both alternatives. The administration every two trials was done to enable generalization to the likely situation in which foregone payoffs are not given all of the time (e.g., when foregone payoffs are revealed by seeing others' choices and outcomes; see discussion section).

Note that although the 1/200 condition actually had a 1 in 200 chance for the –300 outcome we somewhat truncated the variance of this outcome in order to make sure that the results are not derived by participants never seeing the highly negative outcomes. We imposed a 1-3 range for the frequency of –300 occurrences in 400 choices from R. This effectively leaves the average of –300 occurrences in 400 trials as 2 (the actual average was 1.95), but reduces the standard deviation from 1.44 to 0.78. The payoffs series were randomly generated and were different for each participant in a given condition. The forgone payoffs of R were drawn from a separate payoff series than the actual payoffs from R, ensuring that they had no value in predicting the next occurrence of a large loss. However, the payoff series were matched, so that identical payoff series were experienced in the Foregone-1 and Foregone-0 condition (e.g., the same payoff occurred in the *n*'th choice of R for each two participants in the Foregone-0 and 1 conditions). Adding these constraints in the simulation presents no difference from the results depicted in Figure 2.

Recently, it has been found that the effect of experience in experience-based choice tasks is relatively robust to available descriptions of the distributions of the different outcomes (see Yechiam, Barron & Erev, 2005). To examine once again whether experience overwhelms the displayed information, players were given a description of the payoff distributions of the two problems (see Figure 1). It was hypothesized that in both choice problems but in particular in the 1/200 problem, in the Foregone-1 condition there would be an increase the proportion of risky choices from the proportion observed in the condition with no foregone payoffs, and that the differential effect of foregone payoffs would increase with task repetition.

Finally, to examine the heterogeneity in the effect of foregone payoffs we conducted an analysis of individual decision makers. The current model assumes partial weighting of foregone payoffs, implying that the weight of foregone payoffs would be multiplied by a parameter γ that is bounded between 0 and 1. However, an alternative account is that foregone payoffs lead to a gambler's fallacy. For example, when participants receive feedback that they could have lost 300 cents if they chose R, they might feel that they can now confidently switch to choosing the risky option R, which is now "due" for a run of less bad outcomes. If this tendency is consistent, then foregone payoffs are negatively weighted ($\gamma < 0$). To examine individual differences in the response to foregone payoffs, we ran an analysis estimating the model parameters for each participant (for a similar approach, see Yechiam, Busemeyer, Stout & Bechara, in press).

Method

Participants. Eighty students at Indiana University, Bloomington campus (40 males and 40 females), participated in the experiment. Their average age was 20, ranging from 18 to 29. They were paid a sum of \$5 to \$14 for their participation, depending on their success in the experimental task. Participants were randomly assigned to four experimental groups with an equal proportion of males and females in each condition.

Procedure and Apparatus. The experiment took place in the Experimental Spatial Lab at Indiana University. Participants were asked to read the instructions, which were also read out loud. They were encouraged to ask questions. The instruction read as follows: “Your payoff in this experiment will be $18 minus your loses during the experiment. Loses will be accumulated during 400 trials. In each trial you will have to click a button. The payoff for your selection will appear on the button that you selected. You will immediately see a form with two buttons. You can press any of the two buttons in the form (a picture of the form was shown at this point; see Figure 1). The payoff for choosing a button appears below the respective button”.

For the 1/200 Problem participants were instructed: ”In one button there is a probability of 1 in 200 to lose 8 cents and otherwise you lose 2 cent. In the other button there is a probability of 1 in 200 to lose 300 cents and otherwise you lose 1 cent.” For the 1/20 Problem participants were instructed: ”In one button there is a probability of 1 in 200 to lose 8 cents and otherwise you lose 2 cent. In the other button there is a probability of 1 in 20 to lose 300 cents and otherwise you lose 1 cent.” In addition, for both problems the experimenter explained verbally: “This means that in each trial there is a probability of losing a certain amount of money if you press a button, and otherwise you lose a different amount of money. This probability is determined by a random lottery in each trial”.

Participants in the Foregone-1 condition were further instructed as follows:

“In addition, every two rounds you will see information about what is happening in the button that you did not choose. This information will appear on the button that you did not choose”.

All participants were then asked to press the “start the task” button. This presented the game form. The game form included two buttons; the size of each was 3.5 by 6 cubic cm. The actual probabilities and payoffs appeared above each button (see Figure 1). Payoffs were contingent upon the button chosen (S or R) and were calculated independently in each trial per the instructions. The positions of the S and R options (right or left button) were randomized for each player. Two types of feedback immediately followed each choice under all four conditions: (1) The payoff for the choice, which appeared on the selected button until the next button was selected, and (2) an accumulating payoff-counter that was displayed constantly. In addition, in the Foregone-1 condition, players saw the outcome from the button that they did not choose once every two trials starting from trial 1. It appeared on the unselected button until the next button was selected. The foregone payoff appeared regardless of which button was chosen.

Design. The experiment used a 2 × 2 × 4 between and within subjects design with choice problem (1/20 vs. 1/200) and foregone payoffs (Foregone-0 and Foregone-1) as between subject factors and experience (in four blocks of 100 trials) as a within subject factor. Because we used 100 observations per block the sampling distribution of the choice proportion is approximately normal. To stabilize the variances and make the data more compatible with the standard homogeneity of variance assumption of the ANOVA, we conducted the analysis of variance using logit transformations (Logit(p) = ln [p/(1-p)]).

To examine potential heterogeneity in the response to foregone payoffs, the model (described in the appendix) was fitted for each individual in the Foregone-1 condition, allowing for a reverse effect of foregone payoffs (i.e., an effect in the negative direction to the sign of the payoff). For the complete procedure of the estimation, see Yechiam and Busemeyer (2005a). Briefly, The model parameters are fitted to maximize the accuracy (measured as log likelihood) of the `one step ahead' predictions of choices based on the previous outcomes obtained by the participant. The accuracy of the learning model was compared to two baseline models: (1) A random model, (2) a learning model that assumes no effect of foregone payoffs (i.e., the same model with γ set to 0). The statistical test for comparing the fit of the decision model to the baseline models is G^2 (= 2 × log likelihood difference between models) which is a model fit statistic analogous to the chi-square. In addition, because the baseline models have fewer parameters, when we fit parameters, we adjust for the difference in number of parameters. This is achieved by using the Bayesian Information Criterion (BIC; Schwartz, 1978) statistic to compare models. The BIC is a correction that penalizes models with additional parameters: BIC = G^2 - $k \cdot \ln(N)$; where k equals the difference in number parameters and N equals the number of observations. For example, if we have k = 3 (three additional parameter in the learning model compared to the random baseline model) and N = 400, then $3 \cdot \ln(400) \approx 18$. This can be thought of as the deduction from the G^2 of the learning models. Positive values of the BIC statistic indicate that the present learning model performs better than the baseline model, whereas negative values indicate the reverse.

Results and discussion

Figure 3 presents the proportion of selections from option R in each of the four conditions as a function of task repetition. Under both choice problems, in the Foregone-0 condition, following repeated presentations of the task, performers made fewer selections from option R. On average, choices of the risky alternative decreased from 65% in the first block of 50 trials to 40% in the last block.

It appears, though, that in the two foregone payoff conditions the decrease in the choice of the risky alternative was restrained: On average, choices from the risky alternative decreased from 67% in the first block of 50 trials to only 64% in the last block. As predicted, the largest effect of foregone payoff was evident in the 1/200 problem. For this problem, in the Foregone-1 condition, the initial proportion of option R choices was 75%, and the proportion in the final block of trials was only 70%[3].

To examine the statistical significance of this pattern, the results were submitted to a between and within analyses of variance with choice problem (1/20 vs. 1/200) and foregone payoffs (Foregone-0 and Foregone-1) as between subject variables and experience (in four blocks of 100 trials) as a within subject variable. The results showed a main effect of foregone payoffs ($F(1,76) = 4.05$, $p < .05$; $MSE = 7.40$) denoting the increase in the selection from option R in the Foregone-1 condition under both choice problems. In addition, there was a significant interaction between time and foregone payoffs ($F(3,74) = 2.86$, $p < .05$; $MSE = 0.88$), denoting the *increase* of the effect in repeated selections. Finally, there was also a three-way interaction of time, foregone payoffs, and choice problem ($F(3,74) = 2.76$, $p < .05$; $MSE = 0.88$). As expected, the

[3] In addition to examining choice proportion we also examined the actual payoffs of participants. Under Problem 1/200 in the foregone payoff condition, losses were higher (-$9.58) than in the other three conditions (1/200 Foregone-0: -$9.04, 1/20 Foregone-0: -$9.01, 1/20 Foregone-1: -$9.07).

effect of foregone payoffs was stronger in the 1/200 problem, most notably in the second part of the task. In addition to these predicted effects, there was a significant effect of problem ($F(1,76) = 4.05$, $p < .05$; $MSE = 7.40$), with more risky choices in the 1/200 than in the 1/20 problem (consistant with Barron & Erev, 2003; Yechiam & Busemeyer, 2005a).

A post-hoc block by block comparison of the two foregone payoff conditions in Problem 1/200 shows that in the first block of trials there was no significant differences between the Foregone-0 and Foregone-1 conditions ($F(1,38) = 1.67$, $p=.20$). In the last block of trials the effect of foregone payoffs became significant ($F = 5.07$, $p < .05$, $MSE = 3.67$). For an examination of statistical power, *cohen's d* test was used to examine block by block differences. In the first block of 100 trials *d* equaled 0.17 (a small effect) and its value increased to 0.40 (a medium effect size) in the last block. Thus, in problem 1/200, the addition of foregone payoffs increased the proportion of risky selections. Furthermore, the magnitude of the effect was highest following repeated presentations of the task. The post-hoc test for the 1/20 problem showed no significant differences.

Note that paradoxically, the effect of foregone payoffs in the 1/200 Problem occurred even though foregone payoffs increased the average number of times that decision makers saw the –300 outcomes of alternative R. In the Foregone-0 condition the average number of –300 occurrences (in the chosen option) was 1.05 ($SD = 1.1$). In the foregone-1 condition the average number of –300 occurrences (in the chosen or unchosen options) increased to 1.80 ($SD = 1.0$), a significantly larger number of times ($t(38) = 2.31$, $p < .05$). Thus, even though in the Foregone-1 condition participants were exposed more to the highly negative outcome of the risky alternative, they still made more selections from that alternative.

Examination of individual difference. The learning model was fitted for each individual in the Foregone-1 condition. The results show that overall the model fit was much better than the random baseline (Median BIC = 179.5; BIC > 0 for 95% of the participants) and similar to the baseline model that assumes no effect of foregone payoffs (Median BIC = 0; BIC > 0 for 50% of the participants). For the participants in the two choice problems who have observed highly negative foregone outcomes (-30 or -300), the fit compared to the latter model was better for most individuals (Median BIC = 3.6; BIC > 0 for 63% of the participants).

Given that the fit was adequate, we continued to examine the distribution of parameter values in the Foregone-1 condition. Table 1 summarizes the parameter values and figure 4 shows the distribution of the parameter γ, denoting the weight assigned to foregone payoffs in the different choice problems. We subdivided the participants in the 1/200 Problem into the participants who did not see the –300 *foregone* outcome (75%) and those who have seen it (25%). The results in figure 1 show that surprisingly, there were very large individual differences. About 50% of the participants in the 1/20 Problem, and 3 of the 5 participants in the 1/200 Problem who have seen the -300 foregone outcomes, responded to foregone payoffs in a negative direction, implying that if they got a large negative outcome upon selecting alternative R they subsequently made more choices from this alternative.

This suggests that the reinforcement learning model with the prior set parameters explains only some of the variance in the effect of foregone payoffs. Another contributing factor is a gambler's fallacy effect. Indeed, an analysis of the ten individuals with $\gamma < 0$ in the 1/20 problem reveals a median increase of 14% in the proportion of R choices from the ten trials before the –30 foregone outcome to the ten trials afterwards. In contrast, for the ten individuals with positive γ, there was a median decrease of 21% in the proportion of R choices in the ten trials following the –30 foregone outcome. This is consistent with the argument that individuals with $\gamma < 0$ responded to highly negative outcomes from R by making more choices from R.

Finally, we examined whether there are differences in the R choices of the two types of individuals in Problem 1/20, those who partially weight foregone payoffs and those with gambler's fallacy. The average proportions of R choices were very similar in the two groups ($\gamma > 0$: 0.51 R choices, $\gamma < 0$: 0.48 R choices) and a group by block analysis of variance revealed no significant differences, suggesting that both partial weighting and negative weighting of foregone payoffs lead to making more risky choices.

General Discussion

The results of the present experiment support our prediction that when a risky alternative is characterized by rare and highly negative outcomes, this alternative becomes more attractive with the addition of foregone payoffs, and the effect of foregone payoffs on risk taking does not disappear with repeated selections. The present results indicate that the effect of foregone payoffs increases as a function of the degree of rareness of the negative outcome even though the size of this outcome is linearly increased by the same degree.

The present findings could be viewed as paradoxical. Adding descriptions of the probabilities and outcomes of events was found to lead to overweighting of small probability events (see Barron & Erev, 2003; Yechiam, Barron & Erev, 2005). However, increasing the information by adding foregone payoffs leads to more extreme underweighting of small probability events. It was hypothesized that this effect of foregone payoffs occurs because (a) frequent favorable foregone outcomes attract decision-makers to select from the risky alternative, and (b) large negative foregone outcomes are discounted exponentially as a function of experience, and therefore do not impeded the choice of the risky alternative. However, an analysis at the individual level suggested that while for some individuals large negative foregone outcomes are partially weighted and discounted, for others these large negative outcomes are negatively weighted, leading to more choices from the risky alternative. The large negative foregone outcomes are treated as if they were positive outcomes, probably because they are perceived as signals that because the rare event has occurred in the given trial, it will not occur afterwards, and it is safe to choose from the risky alternative. This 'gambler's fallacy' type of behavior most likely reflects misunderstanding of the concept of probability as elucidated by Tverksy and Kahneman (1974). In fact, participants often portray the real world as containing sequential dependence even if given no information of this sort (Budescu & Fischer, 2001).

Empirically, the present results are considered to compliment the findings on the effect of experience on risk taking. They show that at least one regularity of "experience based choice", the tendency to underweight small probability events occurs more strongly and is more enduring in repeated choices with the addition of foregone payoffs. The present results are further supported by a recent study that focused on the choices of student drug abusers compared to non-abusers (Yechiam et al., 2005). It was predicted that drug abusers would be distracted more by foregone payoffs (based on Finn ,2002). The analysis focused on a complex task known as the Iowa Gambling task (Bechara, Damasio, Damasio & Anderson, 1994). The Iowa Gambling task is a four alternative repeated choice task, which initially gives no information about the outcomes from the alternatives, and this has to be learned in repeated selection. Yechiam et al. (2005) used two versions of the task, and the payoffs appear in Table 2. The task was delivered with and without foregone payoffs. The results showed that under the foregone payoff condition, for both payoff versions more choices were made from the two alternatives with a low frequency payoff (B and D). Moreover, for alternative B, which produces constant high positive outcomes and relatively infrequent high negative outcomes, the effect in the drug abusers group did not disappear with repeated experience. The present experiment goes beyond the Yechiam et al. (2005) experiment in several respects. First, the present study demonstrated that this phenomenon occurs for a general student population sample (although drug abusers may certainly be more sensitive to the effect). Second, the present study has elucidated the contributing factors to this phenomenon, in particular the rareness of the negative outcome. Finally, in Yechiam et al. (2005) no descriptions of the alternatives were given prior to the study. The current study has shown that even though decision makers were given accurate descriptions of the probabilities and outcomes, the effect of foregone payoffs still led to large changes in their choices.

As far as the authors know, the only previously published examination of how people combine descriptions and experience is in Yechiam, Barron and Erev (2005) who examined the effect of experience when decision makers were given descriptions of the gambles. They have shown that experience still led to underweighting of small probability events compared to overweighting when participants were given descriptions of the gambles but no experience (the problem they used was a 1/200 problem with the only differences being that the payoffs were in NIS, and the large negative outcome was 200 Agora). The current study shows that foregone payoffs overwhelm the effect of descriptions in a similar manner. It might be that the initial description is entered as an initial propensity for choosing the alternatives (see Barron, Stack & Leider, 2005). However, in situations such as the present problems where the common outcome from the risky alternative is superior to the outcome from the safe alternative, this experience is sufficient to overcome the initial tendency to avoid risky alternatives. Because in real-world situations, there are often both descriptions and experience, this is certainly a topic worthy of perusal in future studies.

<u>Potential limitations</u>

Note that the present reinforcement-learning model predicted the ranking of the alternatives. Specifically, it predicted (a) the direction of the effect of foregone payoffs and its increase in repeated choices, (b) a larger effect of foregone payoffs in Problem 1/200. However, the model under-predicted the effect of the choice problem in the Foregone-0 condition. The overall MSD of the model was only 0.02. Better calibration can potentially be achieved by using more complex choice models (such as Erev and

Barron's (in press) RELACS model) which assume that decision makers weigh-in a cognitive strategy of "avoid loss" in addition to the so called "stage game strategies". Yet our purpose was to demonstrate that even a simple reinforcement- learning model with exponential discounting of old outcomes and partial weighting of foregone payoffs can capture the differential effects of foregone payoffs under varying degrees of rarity of negative payoffs.

A second limitation of the present study is in its generality. The present effects of foregone payoffs were predicted and studied in small repeated decisions that occur in the loss domain, after being given an initial amount of money (i.e., "house money"). When the rare events are in the gain domain, according to the same reinforcement-learning model, foregone payoffs are expected to have an opposite effect, and lead to *overweighting* small probability events. For example, consider a decision maker who faces problem 5:

Problem 5.	S (Safe)	Win 8 cents with a probability of 0.005 (1 in 200)
		Win 2 cents otherwise
	R (Risky)	Win 300 cents with a probability of 0.005 (1 in 200)
		Win 1 cent otherwise

It is likely that in this case the foregone payoff increase the selection from option S, as the most frequent outcome of S (win of 2 cents) is higher than the most frequent outcome of R (win of 1 cent). Thus, the present finding that foregone payoffs increase risk taking is assumed to be limited to a small probability loss. Yet we believe that the effect in the loss domain has some important implications, and these are described next.

On foregone payoffs and social learning

In real-world settings foregone payoffs can be presented by such means as seeing the outcomes of others' selections (see Zeelenberg & Pieters, 2004). This effect of foregone payoffs is often labeled under social or vicarious learning, and can be a mechanism for preventing phobias and neophobias (avoidance of new situations, objects, or environments), which very often are cases of being afraid from things that very rarely lead to harm (Kleinknecht, 1982; see also Jones & Menzies, 2000).

Seeing the outcomes of others' behavior was found to be a mechanism by which animals overcome their fear of eating poisonous food. Animal studies found that when an animal which has been poisoned in the past by a certain food sees that animals in the pack continue to eat similar food, then food aversion learning ameliorates after a smaller number of feeding sessions then when no animals are seen eating the food (Galef, 1987; Provenza & Burritt, 1991; Yoerg, 1991). Animals are also sensitive to harmful outcomes occurring to other animals upon selecting a certain food. If they had previous food aversion to a certain substance, rats tend to avoid this substance more after smelling it on a dead conspecific (Hishimura, 1998).

The present study suggests that foregone payoff information is more effective for reducing sensitivity to extreme low probability but severe risks as compared to higher probability but lower magnitude risks. Understanding this effect of foregone payoffs may be useful for predicting when social learning is likely to lead to long-term changes in behavior. For example, assume that a banana food parasite is more rare but leads to severe outcomes (e.g., prolonged illness) while a mango food parasite is more frequent but leads to a far less severe outcomes (e.g., nausea and vomiting). The current model would predict that risky behavior on the part of conspecifics would have more long-term effects in the case of the bananas. In humans, a similar argument can be made for the long-term effect of peer behavior. For example, it is predicted that certain deviant peer behaviors, such as shoplifting and vandalism, would have a long-lasting effect on adolescents because these behaviors lead to immediate gratification along with rare but severe outcomes (i.e., the chance of getting caught). In contrast, the effect of peer behavior is predicted to be less persistent for risky behaviors with more frequent but less negative outcomes, such as stealing from family members, being rude to peers, etc. Likewise, the present study indicates that foregone payoffs may be helpful in adapting to extreme low probability but severe risks, such as tornadoes, terrorist attacks, etc. For example, the terrorist attacks in the Jerusalem (in the Al-Aqsa Intifada that started in 9.2000) focused on leisure establishments, such as restaurants and shops. A manipulation that includes foregone payoffs might be useful in increasing the proportion of Jerusalem residents who visit these establishments. These selective effects of foregone payoffs are an interesting topic for future investigations.

Appendix 1: The Reinforcement-Learning model

1. Updating of utilities. The current problems involve only losses. Accordingly, we used a correction for dealing with negative expectancies without adding additional parameters (Bereby-Meyer and Erev, 1998; see also Yechiam & Busemeyer, 2005b). The correction, called a Low Reference Point (LRP) solution, forces the expectancies to be positive by deducting the worst possible outcome (in a given task) from all payoffs, as follows: $u(t) = Loss - Maxloss$; where *Loss* is the loss experienced in trial *t, Maxloss* is the maximal loss experienced until trial *t*, and $u(t)$ is the corrected utility. This correction substantially improves the fit for the current problems (Average BIC = 63.9; BIC > 0 for 75% of the participants).

2. Updating of expectancies. Decision makers are assumed to form expectancies for each option, which represent the anticipated consequences of choosing an option. When an option is chosen, the expectancy E_j for option *j* is updated as a function of its previous value (which reflects the past experience), as well as on the basis of a newly experienced payoff *u* on the current trial *t*, as follows:

$$E_{j(t)}= E_{j(t-1)} + \phi\cdot\delta_{j(t)}\cdot[\, u_j(t) - E_j(t-1)] \tag{1}$$

where $\delta_j(t)$ equals 1 if payoff information from option j is presented on trial *t*, and 0 otherwise. This so called “Delta” learning rule was applied to learning in repeated play of individual decision tasks by Busemeyer and Myung (1992), Busemeyer and Stout (2002), and Yechiam and Busemeyer (2005a,b). It has also been applied to repeated play games by Sarin and Vahid (1999) and Erev and Barron (in press). The degree of discounting is determined by the paramater ϕ. High values of ϕ indicate more discounting of old expectancies and more weightings of recent outcomes as opposed to past outcomes. The value of ϕ is limited between 0 and 1, denoting an exponential decrease in the outcome from an alternative as a function of the number of times a particular alternative was chosen.

3. Decision rule for choosing between alternatives. The probability of choosing option 1 to k is determined ratio-of-strength choice rule (Luce, 1959):

$$\Pr[G_j(t+1)] = \frac{e^{\theta(t)\cdot E_j(t)}}{\sum_k e^{\theta(t)\cdot E_k(t)}} \tag{2}$$

Consistency is assumed to increase with experience, reflecting greater reliance of choice on one's expectancies. This is formalized by a power function for the sensitivity change over trials: $\theta(t) = (t/10)^c$, where c is the response sensitivity parameter. When the value of the response sensitivity parameter is high, choices converge towards the alternative with the maximum expectancy. When the value of c is low, choices become inconsistent, random, and independent of the expectancies over time. The value of c is set between -5 and 5, permitting the full range between a deterministic and a random choice.

4. Attention to foregone payoffs. When foregone payoffs are given, formula 3 replaces formula 1:

$$E_j(t)= E_j(t\text{-}1) + \phi\cdot[\delta_j(t) + \gamma / (j\text{-}1)\cdot(1\text{-}\delta_j(t))]\cdot[u_j(t) – E_j(t\text{-}1)] \tag{3}$$

The parameter γ denotes the weight assigned to payoff feedback from non-chosen options, and $u_j(t)$ is the payoff for option j. A γ value of 0 implies no weight, and the model reduces to the same as in formula 1. A value of 1 implies full weight to foregone payoffs in a binary task. Values between 0 and 1 denote partial weighting of foregone payoffs compared to actual (experienced) payoffs). Values above 1 denote increased weighting of foregone payoffs compared to actual payoffs, and values below 0 denote an effect in the reverse direction (or a gambler's fallacy effect) of foregone payoff. The weighting of foregone payoffs from a single alternative is assumed to decrease as a function of the number of alternatives (e.g., in the Iowa Gambling task compared to the current task). However, removing this constraint replicates the current predictions although with a small effect size. Note that as in the experiment, foregone payoffs were presented once every two trials.

5. Parameter values. A parameter estimation on a different choice task (The Iowa Gambling task; Bechara et al., 1994) yielded the following values: $\phi = 0.095$, $c = 0.606$, $\gamma = 0.32$.

References

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.

Barron, G., Stack, J., & Leider, S. (2005). When to warn? The effect of experience on a warnings' impact. *Presented in the Biennial Conference on Subjective Probability, Utility, and Decision Making (SPUDM) 20*. Stockholm, Sweden.

Bechara, A., Damasio, A.R., Damasio, H., & Anderson, S.W. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition*, **50**, 7-15.

Bereby-Meyer, Y., & Erev, I. (1998). On learning to become a successful loser: A comparison of alternative abstractions of learning processes in the loss domain. Journal of Mathematical Psychology, **42**, 266-286.

Budescu, D.V., & Fischer, I. (2001). The same but different: An empirical examination of the reducibility principle. *Journal of Behavioral Decision Making*, **14**, 187-206.

Busemeyer, J.R., & Myung, I.J. (1992). An adaptive approach to human decision-making: learning theory, decision theory, and human performance. *Journal of Experimental Psychology: General*, **121**, 177-194.

Busmeyer, J.R., & Stout, J.C. (2002). A contribution of cognitive decision models to clinical assessment: Decomposing performance on the Bechara gambling task. *Psychological Assessment*, **14**, 253-262.

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

Charness, G., & Grosskopf, B. (2004). What makes cheap talk effective? Experimental evidence. *Economics Letters*, **83**, 383–389.

Erev, I., & Barron, G. (In press). On Adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review*.

Erev, I., Haruvy, E., & Yechiam, E. (2003). A strategic problem in small decisions: Implications to skill learning. Technion Working Paper. See: http://www.indiana.edu/~eypapers/small_dec.pdf

Erev, I., & Rapoport, A. (1998). Coordination, 'magic,' and reinforcement learning in a market entry game. *Games and Economic Behavior*, **23**, 146-175.

Estes, W.K. (1976a). The cognitive side of probability learning. *Psychological Review*, **83**, 37-64.

Estes, W.K. (1976b). Some functions of memory in probability learning and choice behavior. *Psychology of Learning & Motivation*, **10**, 1-45.

Finn, P.R. (2002). Motivation, working memory, and decision making: A cognitive motivational theory of vulnerability to alcoholism. *Behavioral and Cognitive Neuroscience Reviews*, **1**, 181-203.

Galef, B.G. (1987). Social influences on the identification of toxic foods by Norway rats. *Animal Learning and Behavior*, **15**, 327-332.

Gluck, M.A., & Bower, G.H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, **128**, 309-331.

Grosskopf, B. (2003). Reinforcement and directional learning in the ultimatum game with responder competition. *Experimental Economics*, **6**, 141-158.

Grosskopf, B., Erev, I., & Yechiam, E. (2005). Foregone with the wind. Manuscript submitted for publication. See http://www-1.gsb.columbia.edu/faculty/ierev/foregone4.pdf

Haruvy, E. & Erev, I. (2002). Interpreting parameters in learning models. In R. Zwick and A. Rapoport A (Eds.), *Experimental business research*. Kluwer Academic Publishers.

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**, 534-539.

Hishimura, Y. (1998). Food choice in rats (Rattus norvegicus): The effect of exposure to a poisoned conspecific. *Japanese Psychological Research*, **50**, 172-177.

Jones, M.K., & Menzies, R.G. (2000). Danger expectancies, self-efficacy and insight in spider phobia. *Behavior Research and Therapy*, **38**, 585-600.

Josephs, R. A., Larrick, R., Steele, C. M., & Nisbett, R. M. (1992). Self-esteem and risk aversion in decision-making. *Journal of Personality and Social Psychology*, *62*, 26-37.

Luce, R. D., 1959. Individual Choice Behavior. Wiley, New York.

Noble, J., Todd, P.M., & Tuci, E. (2001). Explaining social learning of food preferences without aversions: An evolutionary simulation model of Norway rats. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, **268**, 141-149.

Provenza, F.D., & Burritt, E.A. (1991). Socially induced diet preference ameliorates conditioned food aversion in lambs. *Applied Animal Behavior Science*, **31**, 229-236.

Sarin, R., & Vahid, F. (1999). Payoff assessments without probabilities: A simple dynamic model of choice. *Games and Economic Behavior*, **28**, 294-309.

Sutton, R.S., & Barto, A.G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science,* **185**, 1124-1131.

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.

Yechiam, E., Barron, G., & Erev, I. (2005). The role of personal experience in contributing to different patterns of response to rare terrorist attacks. *Journal of Conflict Resolution*, **49**, 430-439.

Yechiam. E., & Busemeyer, J.R. (2005a). Comparison of basic assumptions embedded in learning models for experience based decision-making. *Psychonomic Bulletin and Review*, **12**, 387-402.

Yechiam. E., & Busemeyer, J.R. (2005b). Evaluating parameter consistency in learning models. Manuscript submitted for publication.

Yechiam, E., Busemeyer, J.R, Stout, J.C, and Bechara, A. (In press) Using cognitive models to map relations between neuropsychological disorders and human decision making deficits. *Psychological Science*.

Yechiam, E., Erev, I., & Gopher, D. (2001). On the potential value and limitations of emphasis change and other exploration-enhancing training methods. *Journal of Experimental Psychology: Applied*, **7**, 277-285.

Yechiam, E., Stout, J.C, Busemeyer, J.R., Rock, S.L., & Finn, P.R. (2005). Individual differences in the response to forgone payoffs: An examination of high functioning drug abusers. *Journal of Behavioral Decision Making*, **18**, 97-110.

Yoerg, S.I. (1991). Social feeding reverses learned flavor aversions in spotted hyenas (Crocuta crocuta). *Journal of Comparative Physiology*, **105**, 185-189.

Zeelenberg. M. (1999). Anticipated regret, expected feedback and behavioral decision making. *Journal of Behavioral Decision Making*, *12*, 93-106.

Zeelenberg, M., & Pieters, R.G.M. (2004). Consequences of regret aversion in real life: The case of the Dutch postcode lottery. *Organizational Behavior and Human Decision Processes*, **93**, 155-168.

Author biography

Eldad Yechiam is a Senior Lecturer at the Faculty of Industrial Engineering and Management at the Technion - Israel Institute of Technology. His research focuses on the implication of basic studies of learning and decision-making for understanding individual differences in complex task performance.

Jerome Busemeyer is Professor of Cognitive Science at Indiana University. His main research interests center on cognitive models of decision making and learning.

Table 1: Medians and standard deviations (in parenthesis) of the estimated parameters for the two choice problems in the Foregone-1 condition.

	Parameter		
Problem	ϕ	c	γ
1/200	0.10 (0.38)	0.74 (0.77)	0.03 (0.67)
1/20	0.05 (0.34)	-0.05 (2.12)	-0.07 (0.65)

Table 2: The payoffs in the Iowa Gambling task in Yechiam, Stout et al. (2005). In the High payoff version the gains and losses from options A and B were increased by 1.5.

	Low payoff version			High Payoff version		
Option	Wins	Losses	Net	Wins	Losses	Net
A	$1.00 for sure	.5 to lose $2.50	-$0.25	$1.50 for sure	.5 to lose $3.75	-$0.375
B	$1.00 for sure	.5 to lose $2.50	-$0. 25	$1.50 for sure	.1 to lose $18.75	-$0.375
C	$0.50 for sure	.5 to lose $0.50	$0.25	$0.50 for sure	.5 to lose $0.50	$0.25
D	$0.50 for sure	.1 to lose $2.50	$0.25	$0.50 for sure	.1 to lose $2.50	$0.25

Figure 1. A screen capture of the experimental task in Experiment 1.

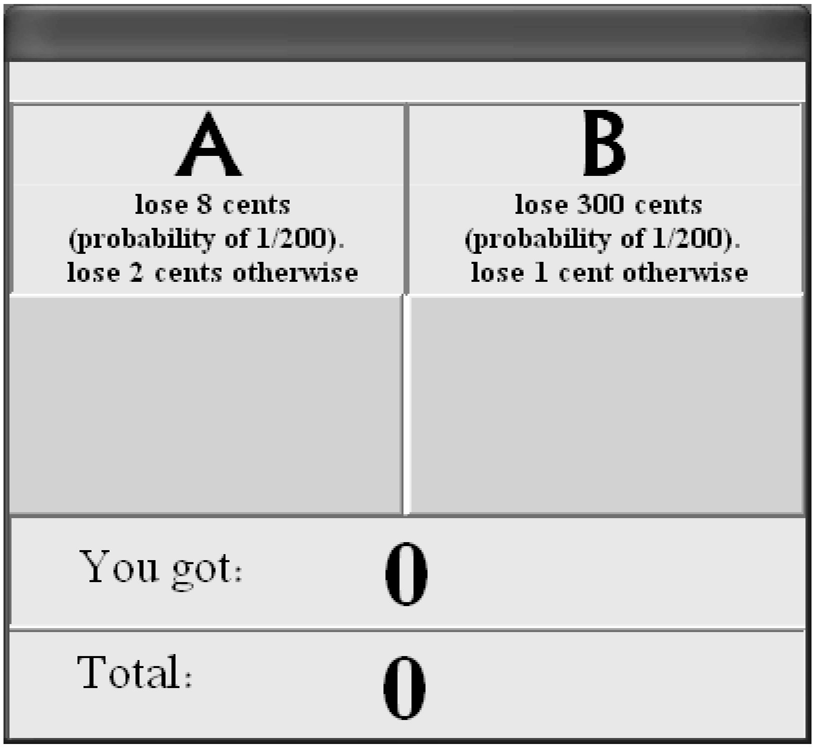

Figure 2. Predicted proportion of selections from option R (the risky alternative) as a function of time (8 blocks of 50 trials). Comparison of Problems 1/200 and 1/20 with foregone payoffs (denoted by F) and without foregone payoff.

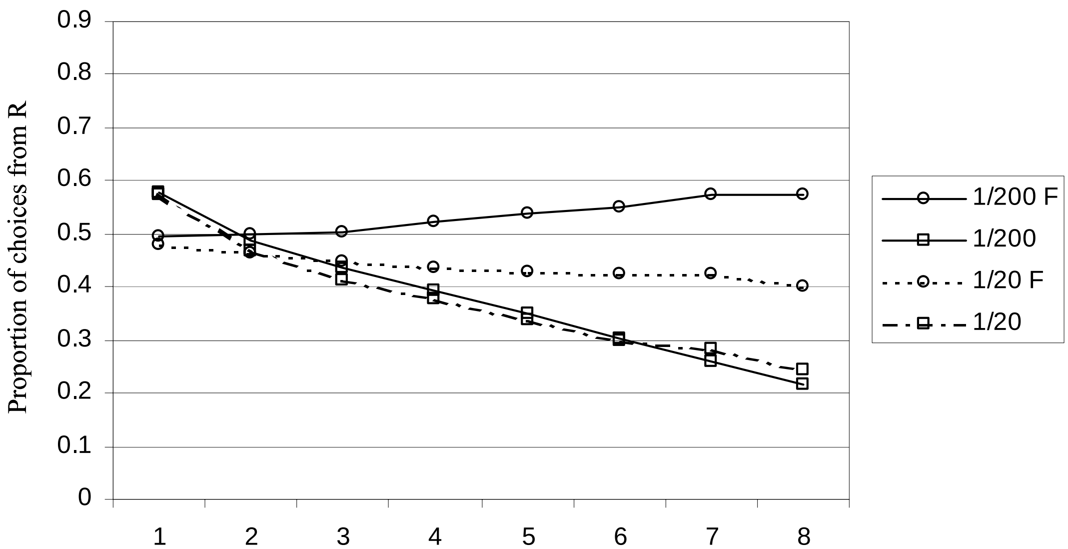

Figure 3. Experiment 1 results: Proportion of selections from option R (the risky alternative) as a function of time (8 blocks of 50 trials) in four experimental conditions: Comparison of Problems 1/200 and 1/20 in the Foregone-1 (denoted by F) and Foregone-0 conditions.

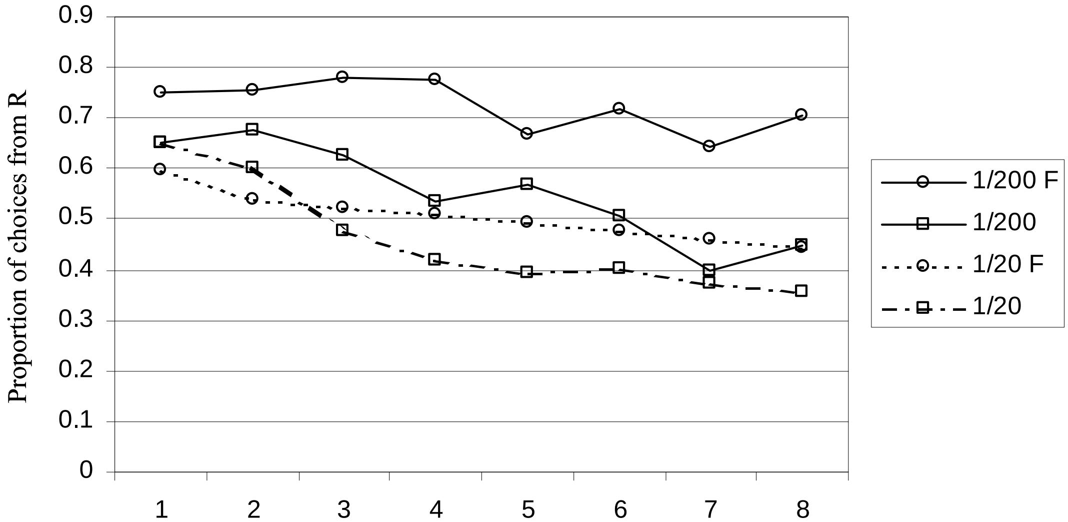

Figure 4. The distribution of the parameter γ, denoting the weight assigned to the foregone payoff feedback, in the two choice problems with foregone payoffs (Foregone-1 condition). The participants in Problem 1/200 are divided into those who viewed the large negative foregone payoff (25%) and those who did not (75%).