Overcoming the Winner's Curse: An Adaptive Learning Perspective

Yoella Bereby-Meyer

Ben Gurion University of the Negev

Brit Grosskopf

Texas A&M University

Friday, February 16, 2007

Running Head: Winner's curse

Mailing addresses

Yoella Bereby-Meyer

Dept. of Education

Ben Gurion University

Beer Sheva 84105, Israel

yoella@bgumail.bgu.ac.il

Brit Grosskopf

Dept. of Economics

Texas A&M University, TAMU 4228

College Station, TX 77843-4228

bgrosskopf@econmail.tamu.edu

Abstract

The winner's curse phenomenon refers to the fact that the winner in a common value auction, in order to actually win the auction, is likely to have overestimated the item's value and consequently is likely to gain less than expected and may even lose (i.e., is said to be “cursed”). Past research, using the “Acquiring a Company” task, has shown that people do not overcome this bias even after they receive extensive feedback. We suggest that the persistence of the winner’s curse is due to a combination of two factors: variability in the environment that leads to ambiguous feedback (i.e., choices and outcomes are only partially correlated) and the tendency of decision makers to learn adaptively. We show in an experiment that by reducing the variance in the feedback, performance can be significantly improved.

Keywords: Winner’s Curse, Asymmetric information, Individual learning, Adaptive behavior.

Overcoming the Winner's Curse: An Adaptive Learning Perspective

The literature in psychology and in economics over the past 50 years has shown that human judgment and decision making deviates at times from rationality (e.g., Allais, 1953; Bazerman, Curhan, Moore, & Valley, 2000; Kahneman & Tversky, 1979; Thaler, 1987; Tversky & Kahneman, 1974). Biases affect all areas of life, including medicine, law, politics, and economics. An important question regarding these biases is whether experience helps people to overcome them.

Past research has identified a phenomenon that is very difficult to avoid, namely the winner's curse. It refers to the fact that the winner in a common value auction, in order to actually win the auction, is likely to have overestimated the item's value and consequently is likely to gain less than expected and may even lose (i.e., is said to be "cursed").

In common value auctions, the value of the auctioned item is the same for all bidders, such as in the case of drilling rights for oil companies. Because it is difficult to estimate the item's exact value, the estimates will vary substantially. In most cases the winning bid will be made by the bidder with the highest estimate of the item's value. Since high estimates are usually not based on more information, the winner will probably overpay for the item. If bidders do not take this *adverse selection* problem into account when deciding on their bids, they will submit winning bids that might produce below normal or even negative profits. Indeed, research has shown that "winning bidders" often find that they have overpaid for the acquired commodities (e.g., Kagel & Levin, 1986).

The winner's curse can also be observed in bargaining under asymmetric information. Take the example of a tourist who arrives at a market in a foreign country and wants to buy a rug. She offers a price for the rug and the seller immediately accepts the offer. Should the buyer learn something from this immediate acceptance? The answer is yes. However, Samuelson and Bazerman (1985) found that negotiators in bilateral bargaining situations with asymmetric information (i.e., where one side has more information than the other side) fail to consider the informational content of their opponent's actions, i.e., they ignore the *selective acceptance* of offers by the more informed side. As a result, they develop inferior bidding strategies that can lead to negative profits (i.e., the winner's curse).

Past research has shown that, even when people have extensive experience with the problem and receive feedback on their performance, they fail to overcome the winner's curse (e.g., Ball, Bazerman, & Carroll, 1991; Foreman & Murnighan, 1996; Selten, Abbink, & Cox, 2005). Little is known about the underlying reasons for this slow learning process. The current paper aims to broaden our knowledge of this topic.

We examine the learning process with a well-researched winner's curse paradigm known as the "Acquiring a Company" task (Samuelson & Bazerman, 1985; a typical set of instructions for this task is given in the Appendix). In the "Acquiring a Company" task, one company considers making an offer to buy another company. The target company's value, v, is known to the seller. However, the buyer only knows that v is uniformly distributed between 0 and 100. If the buyer purchases the company, its value to the buyer will be $1.5 \cdot v$ (i.e., its value will increase by 50%).[1] In these circumstances, what final offer should the buyer make for the company?

Since the seller knows the value of the company, the seller tends to accept offers that are equal or higher than v. This *selective* acceptance of offers leads to a negative expected profit for accepted bids. This can be demonstrated as follows: Assume that a buyer offers $\$b$ for the company. Then the seller will only sell the company, i.e., accept the offer, if the company is worth less than $\$b$. Since v is uniformly distributed with 0 being the lowest possible realization, the expected value of the acquired company is $b/2$. Now, given that the company is worth 50% more to the buyer than it is to the seller, and the buyer needs to pay the offer of b, the overall expected profit is equal to 1.5 times the expected value of the company minus its cost. This is equal to $1.5 \cdot b/2 - b = \frac{3}{4} b - b = -\frac{1}{4} b$. Thus, any offer greater than zero has a negative expected profit and is therefore not worthwhile.

Yet, participants consistently offer amounts that are significantly greater than zero. Most offers lie between the expected value of the company (50) and the ex-ante expected value of the company to the buyer (75) (Samuelson & Bazerman, 1985). In an attempt to understand this finding, Carroll, Bazerman, and Maury (1988) asked participants to think aloud while performing the task. They found that participants simplified their decision task by ignoring the selective acceptance of the seller, thereby treating the problem as if the seller had the same limited information that they had. Carroll et al. (1988) suggested that this simplification of the decision task is a special case of a more general tendency of individuals to make simplifying and potentially biased assumptions when faced with a task that requires incorporating knowledge about future contingent events. Accordingly they showed that performance did not improve when changing the cover story from a bidding scenario to an individual decision making task.

One explanation for participants' inability to take the perspective of others and their resulting choices into account could be the participants' inexperience with the task. It is well accepted that optimal behavior is not observed instantaneously, but rather is likely to evolve through the process of learning and adjustment (Kagel, 1995). Ball, Bazerman, and Carroll (1991) therefore gave participants 20 trials of the “Acquiring a Company” task with financial incentives and full feedback. Participants were informed about the “true” value of the company after each trial, whether their offer was accepted, and how much money they gained or lost. Only 5 out of 69 participants in the experiment lowered their offers over the 20 trials to a value close to zero; the rest of the participants demonstrated almost no learning, i.e., no tendency to lower their bids. In a second experiment the researchers allowed participants to reverse roles in an attempt to make the asymmetry of information more salient to the buyers. For the group that switched roles, the number of learners (defined as participants who bid zero from a certain trial until the end of the experiment) increased from 9 percent to 37 percent. However, for those not defined as learners, only a small reduction in the mean offer was found. Foreman and Murnighan (1996) tried to improve learning in the “Acquiring a Company” task by providing participants with opportunities for both experiential and observational learning. MBA students engaged in 4 repetitions of auctions and the “Acquiring a Company” task over 4 weeks with ample time to reflect and think about the tasks. In addition, the investigators gave students considerable information about bids and offers of other participants and their outcomes. However, neither information, nor experience, nor time, nor a variation on the endowment helped them learn to avoid the curse.[2]

We suggest that one factor that may account for the slow learning in the “Acquiring a Company” task is the noise in the feedback, i.e., the variance in the payoff. By variance in the payoff we refer to two sources of variance: 1.Variance in the probability of winning or losing (we will refer to it from now on as “G/L variance”); and 2. Variance in the amount of payoffs (we will refer to it form now on as “AP variance”).

The “Acquiring a Company” task used by Ball et al. (1991) has a large “G/L variance.” With a 50% premium to the buyer, there is a 1/3 probability of earning positive payoff and 2/3 probability of earning a negative payoff, conditional on an offer being accepted. This can be demonstrated as follows: For an offer of *b,* every company with a value less than *b* will be sold. Since the value of the company to the buyer increases by 50%, the investment will be recouped if the value is exactly 2/3 *b,* since $2/3b \cdot 3/2 = b$. Values smaller than 2/3 *b* will result in a loss and values greater than 2/3 *b* will result in a gain. Thus there is a 1/3 chance of a gain and a 2/3 chance of a loss. The sign of the payoffs (positive, negative) can be modeled as a binomial distribution. For *N* observations and a probability of success of *P*, the variance is equal to:

$$V_{G/L} = N \cdot P \cdot (1-P). \qquad (1)$$

For the current task this $V_{G/L} = 100 \cdot 0.33 \cdot 0.67 = 22.11$. Note that the “G/L variance” does not depend on the value of the bid.

In the “Acquiring a Company” task there is also variance in the amount of payoffs. The “AP variance” for a bid of *b* and a realized value to the buyer of $1.5 \cdot v$ is the variance of a uniform distribution between the minimum possible payoff and the maximum possible payoff. The “AP variance” equals:

$$V_{AP} = \frac{((\min - b) - (1.5 \cdot b - b))^2}{12} = \frac{(\min - 1.5b)^2}{12}, \qquad (2)$$

where min refers to the minimum realized value and equals "0" in the current task.

Note that the “AP variance” depends on the value of the bid.

These sources of variability result in choices and outcomes that are only partially correlated. That means that sub-optimal choices may have high positive payoffs.

Principles of training design suggest that feedback should be frequent and specific (Baldwin & Ford, 1988). If feedback is ambiguous, non-diagnostic, misleading, or poorly managed, it may fail to be corrective, and can even reinforce inappropriate behavior (Einhorn & Hogarth, 1978). According to Tversky and Kahneman (1986), feedback should be accurate and immediate for responsive learning. They postulate that learning is likely to be ineffective if feedback is delayed, if the environment is variable, or if information about foregone payoffs is missing. Recently it has been shown that players’ ability to cooperate in the repeated prisoner's dilemma was substantially reduced when the payoffs were noisy even though players could monitor one another's past actions perfectly (Bereby-Meyer & Roth, 2006).

In the current research we postulate that the failure to overcome the winner’s curse despite extensive experience is due to a combination of two factors: variability in the environment that leads to ambiguous feedback, and the tendency of decision makers to learn adaptively. According to an adaptive learning approach, choices among strategies are not the result of utility calculations but responses that are learned from experience. A favorable experience is assumed to lead to an increase in the propensity to choose a certain strategy, whereas an unfavorable experience is assumed to lead to a decrease in the propensity to choose a strategy (e.g., Bereby-Meyer & Erev, 1998; Grosskopf, 2003; Erev & Roth, 1998; Erev & Barron, 2005; March, 1996; Roth & Erev, 1995; Selten & Buchta, 1998). Since in the “Acquiring a Company” task sub-optimal choices can lead to good outcomes (i.e. high positive payoffs), participants may be misled as to what the optimal choice is, and consequently experience difficulty in learning to avoid the winner's curse (see Erev, Bereby-Meyer, & Roth, 1999, for a reinforcement learning model that explicitly models AP variance as a determinant of learning speed).

In this paper we present an experiment in which we decrease the payoff variance in the “Acquiring a Company” task by averaging the payoffs that participants receive. In this experiment participants are asked to submit an offer that serves as a bid for each of 10 independent companies. After each trial, they receive as payoff the average gain or loss from the attempt to acquire 10 companies simultaneously. The distribution of these mean payoffs forms a much tighter distribution around the negative expected profit, and therefore the probability that the mean payoff of the ten attempts will be positive is smaller than the probability that a single attempt yields a positive payoff. Specifically, the G/L variance of the sampling distribution equals

$$V_{\mathrm{G/L}} = \frac{P*(1-P)}{n} = \frac{P*(1-P)}{10} \qquad (3)$$

It is clear from the comparison of Equations 1 and 3 that the latter is the smaller variance.

Similarly, given the law of large numbers, the variance in the amounts of the rewards will decrease as well. Specifically, the AP variance of this sampling distribution equals

$$V_{Ap} = \frac{((\min - 1.5b)^2/12)}{n} = \frac{((\min - 1.5b)^2/12)}{10} \qquad (4)$$

We show that while holding the expected profit constant, participants learn to avoid the winner’s curse when the variance in the payoffs decreases.

EXPERIMENT

The experiment includes five conditions. The first condition (**Ball-100**) is our basic control condition. It is an extended replication of Ball et al. (1991) for 100 instead of 20 trials. If learning toward the optimal choice occurs, the larger the number of trials, the better the performance should be, given that an asymptote is not yet reached.

In our two experimental conditions (**Average Full** and **Average Only**), we decrease the payoff variance by aggregating and averaging the payoff that participants receive after each trial. In the Average Full condition, participants receive aggregate feedback as well as individual feedback on the performance of each acquisition attempt. It can be argued that this additional information, and not only the reduction in variance, improves performance in the Average Full condition. In order to control for this possibility, we run the Average Only condition without giving the subjects information about the individual companies and giving only information regarding the average payoffs.

We include two additional conditions in order to address common critiques of the "Acquiring a Company" task. One common criticism of the "Acquiring a Company" task in its standard form is that participants might presume that the experimenter wants them to "do something" in the experiment. Therefore choices that are not zero do not necessarily reflect the inability of participants to learn. However, Selten, Abbink, and Cox (2005) found similar overbidding even when the optimal response was to bid a positive amount. Nevertheless, in order to rule out the potential demand effect, we let participants explicitly choose whether to enter into the bidding procedure or to opt out and receive a known payoff (**Yes-No-First**). The assumption was that presenting two options – one for bidding and the other for not bidding – makes the choice set appear more symmetric.

Another explanation for repeated and consistent overbidding can be that participants submit positive bids because they like to gamble. The outcome of a bid of zero is clear – participants break even. However, bidding a positive amount gives the participant an opportunity to participate in a kind of lottery. Even if this lottery has a negative expected value, it sometimes leads to a positive profit. If so, the problem in the "Acquiring a Company" task is not only that high variance prevents reliable estimates of the expected value from bidding, but that participants potentially derive utility from the thrill of gambling, and consequently, the mean observed bid is higher than expected. After all, the gambling industry thrives despite the fact that most people know that the expected value of any given bid is negative. The fifth condition (**Gamble**) tests for this possibility. In this condition, we simplify the "Acquiring a Company" task while maintaining its original structure.

To summarize, given our assertion that decreasing variance should enhance learning in the "Acquiring a Company" task, we hypothesize:

H1: Average bids in Average Full <Average bids in Ball 100.

H2: Average bids in Average Only<Average bids in Ball 100.

Since previous research has shown that task demand is probably not the reason for the slow learning observed in the "Acquiring a Company", we hypothesize that

H3: Average bids in Ball 100=Average bids in Yes-No First.

Method:

Participants

One hundred and nineteen participants were recruited from the Boston area (21 in Ball-100, 28 in Average Full, 26 in Average only, 20 in Yes-No first, and 24 in the Gamble condition). Their ages ranged between 18 and 60 years, and they were paid according to their performance. Participants received a $10 show-up fee and had the opportunity to earn up to $10 more.

Procedure

The experiment was computerized and programmed in zTree (Fischbacher, in press). It consisted of the following five conditions:[3]

1) Ball-100: Participants faced the "Acquiring a Company" task for 100 trials in the role of buyers. In each trial, participants were asked to determine a price they would like to offer for the company. Participants were informed that the value of the company is distributed uniformly between 0 and 100 and that the seller (represented by a computer program) knows the value of the company, while the buyer does not. Participants were also told that the company would be worth 50 percent more under their management than under the management of the current owner. At the beginning of the experiment, each participant received an endowment of 5,000 points, with each point being worth 0.2 cents. The feedback after each trial included: (1) the "true" (realized) value of the company, (2) the value of the company to the buyer, (3) the buyer's most recent profit/loss, and (4) the accumulated payoff.

2) Average Full: The basic procedure in this condition was the same as in Ball-100, except for the following changes. Instead of buying one company in each of the 100 trials, participants were asked to submit one bid that would serve as an offer for each of 10 independent companies. After each trial, participants received feedback regarding their average payoffs from the attempt to buy the 10 companies. In addition, they received feedback regarding the value of each of the 10 companies and the payoff from each individual deal.

3) Average Only: This condition was similar to the Average Full condition, except for the fact that participants received only the average payoff without information on each individual deal.

4) Yes-No First: The basic procedure in this condition was the same as in the Ball-100 condition, except for the following changes. In this condition, participants were first asked to choose between two options: (1) not to bid and (2) to bid. If participants chose not to bid, they received 50 points for sure; if they chose to bid, then the outcome was added to (or subtracted from) the 50 points. The information that was given after each trial was identical to the information given in the Ball-100 condition.

5) Gamble: In each round of this condition, participants were asked to choose one of 10 letters corresponding to a lottery that was conducted in each trial. Participants received a chart that presented the expected payoff from each choice (see Figure 1).

Insert Figure 1 here

The area distributions within the different bars in the chart correspond to the distributions of possible payoffs (i.e., lose, win or break even) associated with the letter choices. The numbers in the different segments refer to the points a participant could gain or lose depending on the outcome of the lottery.[4] On the right-hand side of the chart, the expected payoff for each possible choice was presented. Letters A through K were multiples of 10 (from 0 to 100) for bids in the "Acquiring a Company" task. At the beginning of the 100 trials, participants received 1,500 points. The exchange rate for each point was 0.7 cents. After each trial, participants were informed of their payoff for the trial and their accumulated payoff. Table 1 gives a summary of the 5 conditions in our experiment.

Insert Table 1 here

Results

Our hypothesis was that the variance in the environment, together with the tendency to learn adaptively, impedes learning. Therefore, we will first characterize learning behavior in the Ball-100 condition. In order to demonstrate that subjects behave adaptively in this condition, we analyzed the sequential dependencies in the data, i.e., the dependence of bids in trial *t+1* on the outcome of trial *t*.

Figure 2 presents the direction of changes in bids at time *t+1* as a function of whether the company was not acquired, acquired under losses, or acquired profitably at time *t* (choices of zero were excluded from the analysis).

Insert Figure 2 here

As can be seen from Figure 2, for companies that were acquired, participants gained money one third of the time and lost money two thirds of the time. An ANOVA, in which the outcome at time *t* (null, loss, gain) was a repeated measure variable and the difference between bid at time *t+1* and the bid at time t was the dependent variable, revealed a significant effect for the outcome at time *t*, $F(2,40) = 24.38$, $p<0.0004$, $MSE=118.499$. As predicted by an adaptive learning process, the outcome at time *t* affected the change in bids at time *t+1*. The highest tendency to decrease bids was found when a company was acquired under losses. A smaller decrease was found when a company was profitably acquired. When a company was not acquired, resulting in a zero payoffs, there was a tendency to increase bids (an LSD post-hoc test revealed a significant difference for all comparisons, $p<0.05$).

This bounded rationality behavior can be explained through a directional learning process of the type suggested by Selten and Buchta (1998). According to this learning process, the decision maker looks at the outcome of the previous round and makes her choice in the following round in a qualitatively better direction, i.e., bids more or less. Obviously, the decision maker might ignore potentially valuable information that is masked by the variability in the underlying payoff structure, and consequently experience difficulties in learning to avoid the winner's curse.

Given the demonstrated adaptive nature of the behavior, we shall now turn to analyzing behavior in our experimental conditions in which we lowered the payoff variance. Figure 3 shows the mean bids for every block of 10 trials for the different conditions.[5] As can be seen in the figure, a reduction in mean bids occurred in the Average conditions and also in the Gamble condition while in the Control and the Yes-No First condition almost no reduction in mean bids occurred over time.

A two-way ANOVA with one between-factor (Condition), and one within-factor (Block: 10 blocks of 10 trials each) as independent variables and with the mean bids as a dependent variable revealed a significant effect for the Condition, $F(4,114)=3.6$, $p<0.008$, $MSE=3420$. In accordance with *H1*, mean bids in the Average Full condition were lower than mean bids in the Ball-100 condition, $F(1,114)=10.04$, $p<0.001$, $MSE=3419.81$. Similarly, and supporting *H2*, mean bids in the Average Only condition were significantly lower than mean bids in the Ball-100 condition, $F(1,114)=5.97$, $p<0.02$, $MSE=3419.81$. As predicted by *H3*, no difference was found between mean bids in Ball-100 and mean bids in Yes-No First condition, making the explanation of task demand as a potential explanation less probable. A significant effect was also found for the Block, $F(9,1026)=9.745$, $p<0.0001$, $MSE=209$, and for the interaction between the Block and the Condition, $F(36,1026)=1.97$, $p<0.0006$, $MSE=209$.

Insert Figure 3 here

A linear contrast on the Block in each condition revealed significant linear trends for the Average Full (F $(1,114)=14.41$, $p<0.002$, $MSE=677.182$), for the Average Only (F $(1,114)=4.99$, $p<0.02$, $MSE=677.182$), and the Gamble ($F(1,114)=25.95$, $p<0.0003$, $MSE=677.182$) conditions. No significant trend was found for the other conditions.

Since it might be the case that very few participants learned without necessarily affecting the mean bid, we assessed the number of learners in each condition. Participants were defined as learners if they bid values of less than 10 from any particular trial on until the end of the experiment and through at least the last three trials. The numbers of learners for the Control, Average Full, Average Only, Yes-No First, and Gamble conditions were 4, 14, 4, 5, and 9, respectively. As can be seen, very few participants can be defined as learners in the different conditions, except for the Average Full condition, for which a higher number of learners were found.

These results show that when the variance in the payoffs was reduced, either in the Average Full or in the Average Only condition, participants learned to overcome the winner's curse and bid less than in the Ball-100 condition. Participants also learned in the Gamble condition. However, even in this condition, in which participants received explicit information regarding the expected value of each choice, participants continued to bid. This behavior suggests that participants obtain utility from bidding. The fact that no difference was found in the mean bids of the Yes-No First condition and Ball-100 condition, supplemented with the finding that the trend was not significant in the Yes-No First condition, suggests that the demand characteristics of the task are probably not a potential explanation for the slow learning process that was observed.

General Discussion

A large body of research on bilateral bargaining under uncertainty has found that, under asymmetric information, decision makers often develop inferior bidding strategies that can result in negative profits that lead to the winner's curse phenomenon (Bazerman &

Samuelson, 1983; Samuelson & Bazerman, 1985). Very intriguing is the observed regularity that even with extensive experience, people do not learn to adjust and overcome this bias (Ball, Bazerman, & Carroll, 1991; Foreman & Murnighan, 1996). We replicated this finding in our Ball-100 condition. To date, little was known about the reasons for this slow adaptation process.

We propose an explanation based on research on adaptive learning models. One possible reason for the slow adjustment may be the variance that characterizes the bargaining situation. Sub-optimal choices in the "Acquiring a Company" task can yield positive and high payoffs. In fact, one third of the time participants gain positive amounts of money while bidding and acquiring a company. However, for feedback to improve learning, it has to help in rejecting erroneous hypotheses (Balzer et al., 1989; Kluger & DeNisi, 1996). But after positive payoffs, it apparently becomes very difficult for participants to reject the hypothesis that it is good to bid. Sequential analysis of the Ball-100 condition shows that participants condition their bids on the outcomes of previous trials, behavior that is characteristic of an adaptive learning process.

For this reason, we decreased the variance by modifying the task in our experimental conditions (Average Full and Average Only). Instead of submitting one bid for one company, participants submitted one bid for 10 companies. The mean payoffs from buying 10 companies offers less noisy feedback than the individual payoff from each bid. In accordance with our hypotheses, we found that, given the same expected profit, mean bids decreased when the variance was decreased. In the current research, the "G/L variance" and the "AP variance" were manipulated simultaneously, and consequently we are not able to pin point which of them or maybe both account for the improvement in learning. Future research should try to disentangle these two types of variances.

Participants' tendency to rely on the outcomes of previous rounds instead of calculating expected utility was observed in financial investment decisions, for which a “myopic loss aversion” was suggested (Benartzi & Thaler, 1995). This sub-optimal behavior was overcome when participants received aggregate feedback and could not change their behavior from trial to trial (Gneezy & Potters, 1997; Thaler, Tversky, Kahneman, & Schwartz, 1997). Similarly, Silverstein, Cross, Brown, and Rachlin (1998) found that asking participants to decide on four moves in advance in a prisoner’s dilemma game helps to achieve a higher degree of cooperation. These findings seem to contradict the recommendation of frequent feedback for an effective learning process (e.g., Baldwin & Ford, 1988; Tversky & Kahneman, 1986). When feedback is noisy, less frequent feedback may be preferable to achieve a more stable estimate of the optimal behavior. This is especially important given the frequently observed adaptive learning tendency of decision makers to adjust their behavior to the outcome of the previous experience.

A remaining question is what kind of learning was achieved in this task. Neale and Northcraft (1990) differentiate between experience and expertise. They define experience as a change in behavior brought about by a mindless, adaptive reinforcement process, with limited transferability to new contexts. By contrast, expertise is the development of a schematic principle that involves conceptual understanding of the problem, i.e., an understanding of what one has experienced and awareness of the decision process that led to failure or success. Similarly, Baron (2000) suggests that learning without understanding was characterized by either lack of transfer of the principles to cases where it applied or by inappropriate transfer to cases where it did not apply.

The current study does not enable us to know what kind of learning took place. We can postulate, given the small number of learners in the Average only condition, that participants probably undergo a simple adaptive learning process that leads to the acquisition of habitual behavioral patterns but does not necessarily result in a deep understanding of the problem. Future research should try to characterize the type of learning process more directly, for example, by giving the participants a new transfer task with a different cover story after completing the 100 trials (i.e., in order to examine positive transfer), or by giving them the same task with a different premium value where it is worthwhile to bid and see if they fall prey to negative transfer of not bidding.

While the current research suggests that learning from experience in the "Acquiring a Company" task is very difficult, recently, some research suggests that having participants compare and contrast their choice within and across problems, might shed light on effective learning mechanisms to improve decision making in the "Acquiring a Company" task and similar problems (Idson et al., 2004)

Lastly, the susceptibility of learning to feedback variance suggests that a claim by economists that biases will be eliminated with experience might be restricted to environments with low amounts of noise. Thus, it is important that experiments that test learning be done in noisy environments that better reflect real-life situations.

References

Allais, P. M. (1953). The behavior of rational man in risky situations - A critique of the axioms and postulates of the American School. *Econometrica*, *21*, 503-546.

Baldwin, T., & Ford, J. K. (1988). Transfer of training: A review and directions for future research. *Personnel Psychology*, *41*, 63-105.

Ball, S., Bazerman, M., & Carroll, J. (1991). An evaluation of earning in the bilateral winner's curse. *Organizational Behavior and Human Decision Processes*, *48*, 1-22.

Baron, J. (2000). *Thinking and deciding* (3rd ed.). New York, NY: Cambridge University Press.

Bazerman, M. H., Curhan, J. R., Moore, D. A., & Valley, K. L. (2000). Negotiation. *Annual Review of Psychology*, *51*, 279-314.

Bazerman, M. H., & Samuelson, W. F. (1983). I won the auction but don't want the prize. *Journal of Conflict Resolution*, *27*, 618-634.

Benartzi, S., & Thaler, R. (1995). Myopic loss aversion and the equity premium puzzle. *Quarterly Journal of Economics*, *110*, 73-92.

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.

Bereby-Meyer, Y., & Roth, A. (2006). The speed of learning in noisy games: Partial reinforcement and the sustainability of cooperation. *American Economic Review, 96*, 1029-1042.

Balzer, W. K., Doherty, M. E., & O'Connor, R. (1989). Effects of cognitive feedback on performance. *Psychological Bulletin*, *3*, 410-433.

Carroll, J. S., Bazerman, M. H., & Maury, R. (1988). Negotiator cognitions: A descriptive approach to negotiators' understanding of their opponents. *Organizational Behavior and Human Decision Processes*, *41*, 352-370.

Einhorn, H. J., & Hogarth, R. M. (1978). Confidence in judgment: Persistence of the illusion of validity. *Psychological Review, 85*, 395-416.

Erev, I., & Roth A. E. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review*, *88*, 848-881.

Erev, I., Bereby-Meyer,Y., & Roth, A. (1999). The effect of adding a constant to all payoffs: Experimental investigation, and implications for reinforcement learning models. *Journal of Economic Behavior and Organization*, *39*, 111-128

Erev, I., & Barron, G. (2005). On Adaptation, Maximization, and Reinforcement Learning Among Cognitive Strategies. *Psychological Review, 112,* 912-931.

Fischbacher, U. (in press), z-Tree - Zurich Toolbox for Readymade Economic Experiments. *Experimental Economics*.

Foreman, P., & Murnighan, J. K. (1996). Learning to avoid the winner's curse. *Organizational Behavior and Human Decision Processes*, *67*, 170-180.

Gneezy, U., & Potters, J. (1997). An experiment on risk taking and evaluation periods. *Quarterly Journal of Economic*s, *112*, 631-645.

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

Idson, L. C., Chugh, D., Bereby-Meyer, Y., Moran, S., Grosskopf, B., & Bazerman, M. (2004). Overcoming focusing failures in competitive environment. *Journal of Behavioral Decision Making*, *17*, 159-172.

Kagel, J. H. (1995). Auctions: A survey of experimental research. In J. H. Kagel & A. E. Roth (Eds). *The handbook of experimental economics* (pp. 501-536). Princeton, NJ; Princeton University Press.

Kagel, J.H., & Levin, D. (1986). The winner's curse and public information in common value auctions. *American Economic Review, 76*, 894-920.

Kagel, J.H., & Levin, D. (2002). *Common value auctions and the winner's curse*, Princeton, NJ; Princeton University Press.

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

Kluger, A. N., & DeNisi, A. (1996). The effects of feedback interventions on performance: A historical review, a meta-analysis, and a preliminary feedback intervention theory. *Psychological Bulletin*, *119*, 254-284.

March, J. G. (1996). Learning to be risk averse. *Psychological Review*, *103*, 309-319.

Neale, M. A., & Northcraft, G. B. (1990). Experience, expertise, and decision bias in negotiation: The role of strategic conceptualization. In B. Sheppard, M. Bazerman, & R. Lewicki (Eds.), *Research on negotiations in organizations* (Vol.2, pp. 55-75). Greenwich, CT.: JAI Press.

Roth, A. E., & Erev, I. (1995). Learning in extensive form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior*, *8*, 164-212.

Samuelson, W. F., & Bazerman, M. H.(1985). Negotiation under the winner's curse. In V. Smith (Ed.) *Research in experimental economics* (Vol. III, pp. 38-105). Greenwich, CT, JAI Press.

Selten, R., Abbink, K., & Cox, R. (2005). Learning direction theory and the Winner's Curse. *Experimental Economics*, *8*, 5 – 20.

Selten, R., & Buchta, J. (1998). Experimental sealed bid first price auctions with directly observed bid functions. In D. Budescu, I. Erev, & R. Zwick (Eds.) *Games and human behavior: Essays in honor of Amnon Rapoport* (pp. 79-104). Mawhaw NJ: Lawrence Erlbaum Associates, Inc.

Silverstein, A., Cross, D., Brown, J., & Rachlin, H. (1998). Prior experience and patterning in a Prisoner's Dilemma game. *Journal of Behavioral Decision Making*, *11*, 123-138.

Thaler, R. (1987). The psychology of choice and assumption of economics. In A. E. Roth (Ed.), *Laboratory experimentation in economics: Six points of view* (pp. 99-130). Cambridge: Cambridge University Press.

Thaler, R. H., Tversky, A., Kahneman, D., & Schwartz, A. (1997). The effect of myopia and loss aversion on risk taking: An experimental test. *Quarterly Journal of Economics*, *112*, 647-661.

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

Tversky, A., & Kahneman, D. (1986). Rational choice and the framing of the decisions. *Journal of Business*, *59*, 251-294.

Weinstock, S. (1958). Acquisition and extinction of a partially reinforced running response at a 24-hour intertrial interval. *Journal of Experimental Psychology*, *56*, 151-158.

Footnotes:

1. This is often justified by assuming better management under the new ownership.

2. A more significant adjustment to overcome the winner's curse over time was found in the context of auctions (for a review see Kagel & Levin, 2002). There are some differences between the Acquiring a Company task and an auction. Since the current study focuses on the Acquiring a Company task, we will not elaborate on this issue.

3. Participants were randomly assigned to conditions 1 and 2. The other conditions were run later. However, participants in all conditions where recruited from the same subject pool and through the web. Therefore there is no reason to believe that participants were different across treatments or were able to communicate outside the lab.

4. The original instructions included a colored version of this chart; losses were represented by red bars, gains by blue bars, and payoffs of zero by yellow bars. Red was chosen to emphasize the negative outcomes.

5. In order to compare the "Gamble" condition to the other conditions, we transformed the letters in that condition to numbers from 0 to 100.

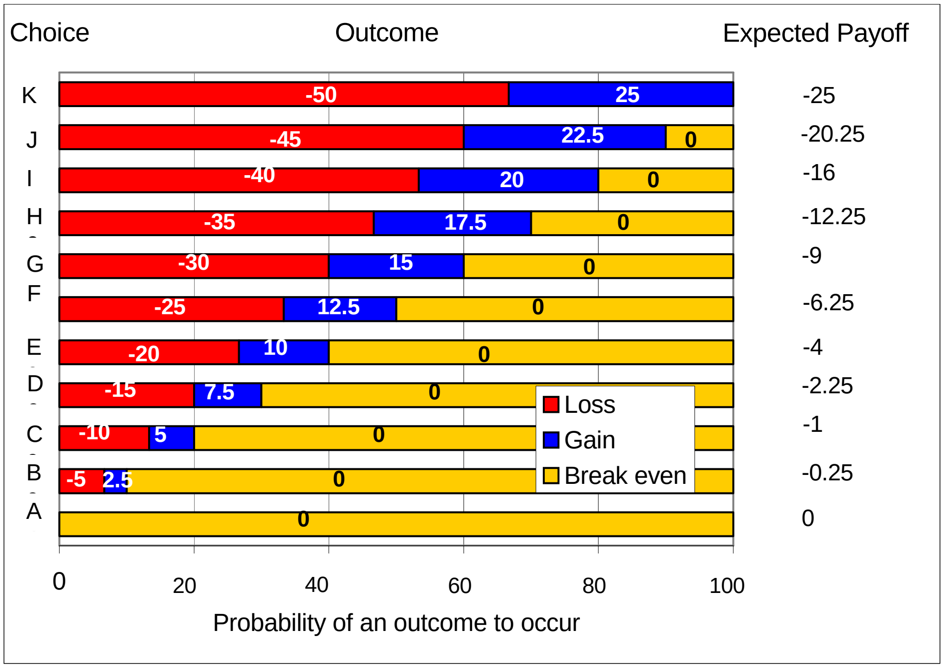

Figure 1**:** Chart that participants were given in the Gamble condition.

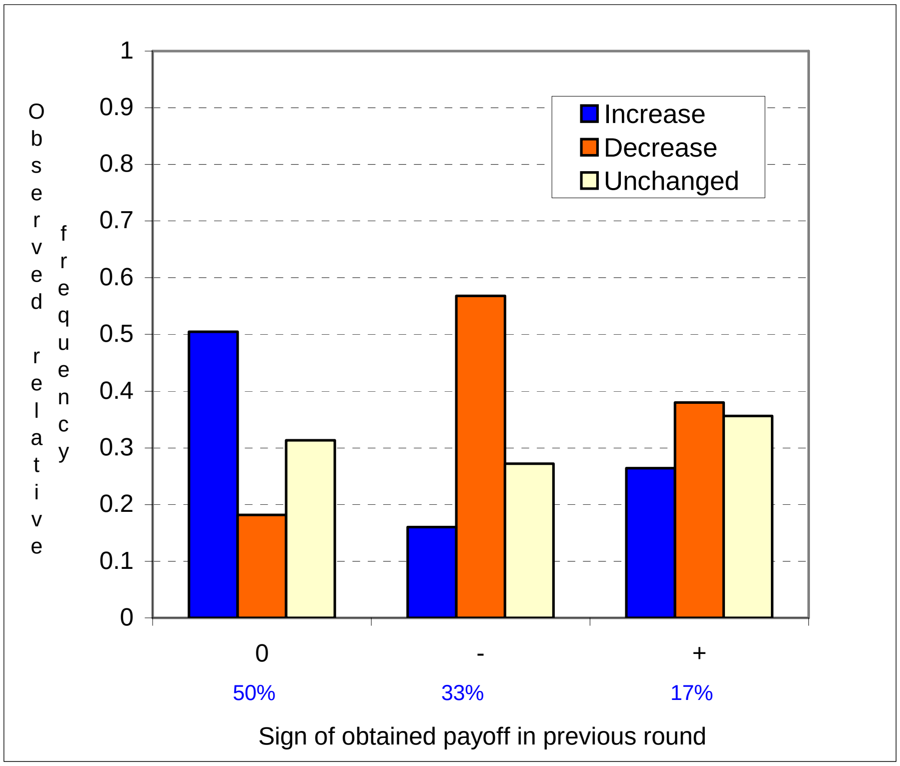

Figure 2**:** The direction of change in bids at time *t+1* as a function of the outcome at time *t*

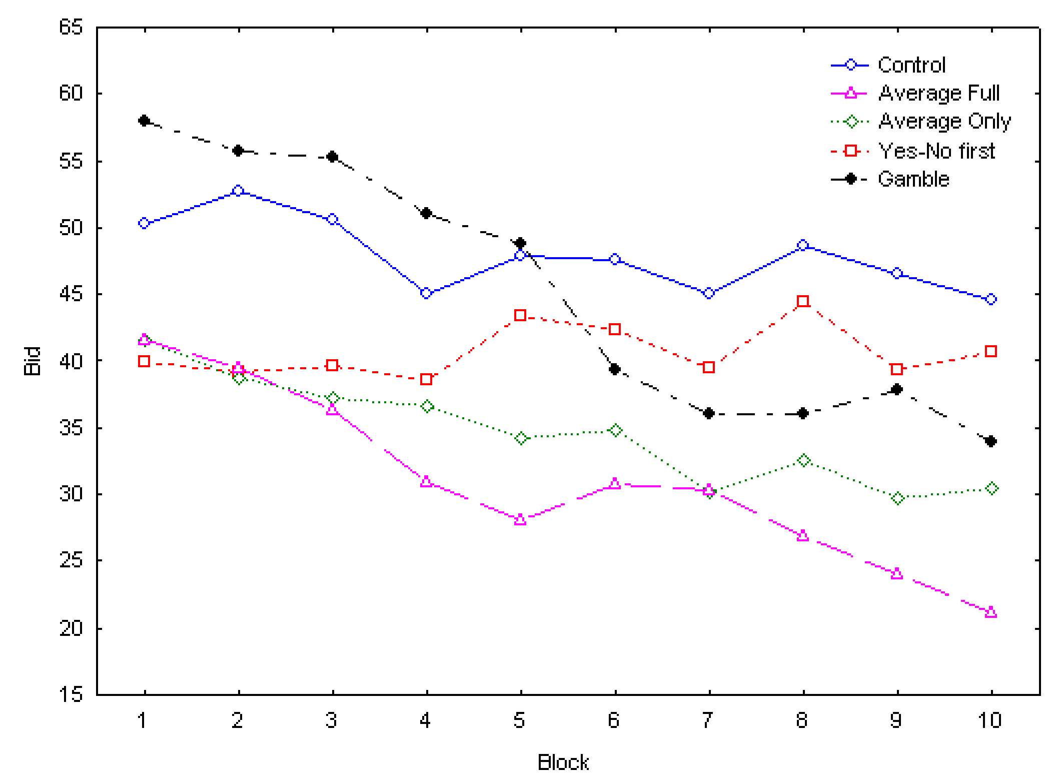

Figure 3**:** Mean bids for every block of 10 trials for all conditions

Table 1: Summary of the conditions (*v*: value to the seller, 1.5*v*: value to the buyer, *π*: earned payoff)

		Number of companies to be acquired	*Feedback*	*Number of participants*
Control Conditions	Ball-100	1	*v*, 1.5*v*, *π*	21
	Yes-No First	1	*v*, 1.5*v*, *π*	20
	Gamble	1	*π*	24
Experimental Conditions	Average Only	10	Average aggreg. *π*	26
	Average Full	10	Individual *v*, 1.5*v*, *π* and Average aggreg. *π*	28

Appendix: Sample set of instructions for the "Acquiring a Company" task

You will represent Company A (the potential acquirer), which is currently considering acquiring Company T (the target) by means of a tender offer. The main complication is this: the value of Company T depends directly on the outcome of a major oil exploration project that is currently undertaking. Indeed, the very viability of Company T depends on the exploration outcome. If the project fails, the company under current management will be worth nothing (0 points per share). But if the project succeeds, the value of the company under current management could be as high as 100 points per share. All share values between 0 points and 100 points are considered equally likely.

By all estimates, Company T will be worth considerably more in the hands of Company A than under current management. In fact, the company will be worth 50 percent more under the management of A than under the management of Company T. If the project fails, the company will be worth 0 points/share under either management. If the exploration project generates a 50 points/share value under current management, the value under Company A will be 75 points/share. Similarly, a 100 points/share value under Company T implies a 150 points/share value under Company A, and so on.

It should be noted that the only possible option to be considered is paying in cash for acquiring 100 percent of Company T's shares. This means that if you acquire Company T, its current management will no longer have any shares in the company, and therefore will not benefit in any way from the increase in its value under your management.

The board of directors of Company A has asked you to determine whether or not to submit an offer for acquiring company T's shares, and if so, what price they should offer for these shares. This offer must be made now, before the outcome of the drilling project is known. Company T will accept any offer from Company A, provided it is at a profitable price for them. It is also clear that Company T will delay its decision to accept or reject your bid until the results of the drilling project are in. Thus you (Company A) will not know the results of the exploration project when submitting your price offer, but Company T will know the results when deciding on your offer. As already explained, Company T is expected to accept any offer by Company A that is greater than the (per share) value of the company under current management, and to reject any offers that are below or equal to this value. Thus, if you offer 60 points/share, for example, Company T will accept if the value under current management is anything less than 60 points.

As the representative of Company A, you are deliberating over price offers in the range of 0 points/share (this is equivalent to making no offer at all) to 150 points/share. Your bid should be expressed in whole points (integer values).

What is your offer?

Yoella Bereby-Meyer is a Senior Lecturer at the Department of Education in Ben-Gurion University. She received a PhD in Industrial Engineering from the Technion-Israel Institute of Technology in 1997. Her areas of research include decision-making, learning, and negotiation.

Brit Grosskopf is an Assistant Professor of economics at Texas A&M University. She holds a Ph.D. in Economics from the Universitat Pompeu Fabra, Barcelona (2000). Her research interests include experimental economics, learning, decision-making, and bounded rationality.