On the Robustness of the Winner's Curse Phenomenon

Brit Grosskopf

Yoella Bereby-Meyer

Max Bazerman

This Draft: May 31, 2006

On the Robustness of the Winner's Curse Phenomenon*

Abstract

We set out to find ways to help decision makers overcome the "winner's curse," a phenomenon commonly observed in asymmetric information bargaining situations, and instead found strong support for its robustness. In a series of manipulations of the "Acquiring a Company Task," we tried to enhance decision makers' cognitive understanding of the task. We did so by presenting them with different parameters of the task, having them compare and contrast these different parameters, giving them full feedback on their history of choices and resulting outcomes, and allowing them to interact with a human opponent instead of a computer program. Much to our surprise, none of these manipulations led to a better understanding of the task. Our results demonstrate and emphasize the robustness of the winner's curse phenomenon.

JEL classification: C78, C91

Keywords: Winner's Curse, Asymmetric Information, Cognitive Learning

* Financial support from the Russell Sage Foundation (# 98-02-01) is gratefully acknowledged. The experimental interface was programmed in zTree, an experimental software authored by Urs Fischbacher (Fischbacher, 1999).

1. Introduction

Abundant evidence demonstrates that decision makers systematically deviate from rational behavior. Deviations from rationality have been found in individual decision-making tasks (e.g. Allais, 1953; Tversky and Kahneman, 1974; Kahneman and Tversky, 1979; Bazerman, 2002), as well as in interactive situations with a social context (e.g. Bazerman, Curhan, Moore and Valley, 2000; Grosskopf and Nagel, 2001). As these biases hinder the quality of decision-making in countless domains, including economics, politics, and medicine, it is important to study ways to overcome them.

The current research focuses on the tendency of decision makers to ignore the decisions of others and the resulting consequences. As a result of this tendency, decision makers commonly develop inferior bidding strategies in bilateral bargaining situations with asymmetric information that lead to negative profits – a phenomenon known as the winner's curse. We examine whether decision makers can learn to overcome this bias if they are given feedback that aims to trigger their cognitive understanding by focusing on the obstacles to learning that were found in prior studies.

In this study we use the "Acquiring a Company Task," (also known as the "Takeover Game") originally studied by Samuelson and Bazerman (1985), and based on Akerlof's (1970) article on the "market for lemons." In the "Acquiring a Company Task," buyers make a take-it-or-leave-it offer to sellers for a company with a value, *v to the seller,* that is uniformly distributed between 0 and 100, and with the value to the buyer being 1.5**v*. However, only the seller knows the exact value of the company at the time bids are submitted. Since one side (the seller) has more information than the other (the buyer), it will only accept offers that are higher or equal to *v*. As a result with these

parameters, the expected value to the seller of any accepted positive bid is negative. The negative expected value is driven by the fact that the seller only accepts offers that are higher than his valuation. The less-informed side (the buyer) should anticipate this *selective* acceptance and its consequences, and thus should not submit any bid at all. Yet previous studies have shown that participants consistently bid amounts that are significantly greater than zero; most bids lie between the expected value of the company (50) and the ex-ante expected value of the company to the buyer (75) (e.g. Bazerman and Samuelson, 1983). An analysis of participants' "think-aloud transcripts" revealed that participants simplified their decision task by ignoring the selective acceptance of the other side that is a consequence of the information asymmetry. Participants seemed to treat the problem as if the seller was not better informed than they were (Carroll, Bazerman and Maury, 1988; Tor and Bazerman, 2003).

An important question regarding biases is whether agents can learn to overcome them. It is well accepted among experimental economists that optimal behavior should not necessarily be expected right away, but is likely to evolve through the process of learning and adjustment (Kagel, 1995). In fact, ample evidence exists to show that behavior frequently, but not necessarily, converges to rationality with experience (e.g. Coursey, Hovis and Schultze, 1984; Smith, 1985; Roth and Erev, 1995). In contrast, participants in the "Acquiring a Company Task" have exhibited a strong persistence in sub-optimal behavior, even when given extended experience with the task. Ball, Bazerman and Carroll (1991) gave participants 20 trials of the "Acquiring a Company Task" with financial incentives and full feedback. After each trial, participants were told the "true" (i.e., realized) value of the company, whether their bid was accepted, and how

much money they gained or lost. Only 5 out of 69 participants in the experiment decreased their bids over the 20 trials to a value close to zero; the rest of the participants demonstrated almost no learning.

Early criticism of the Ball *et al.* study focused on the uniqueness of the right answer (bidding zero), arguing that there exists a demand characteristic to bid, rather than to do nothing. However, Selten, Abbink and Cox (2005) found similar overbidding even when the optimal response was to bid a positive amount and even when participants were given 100 trials to gain experience. Foreman and Murnighan (1996) tried to improve learning in the "Acquiring a Company" task by providing participants with opportunities for both experiential as well as observational learning. In their study, MBA students engaged in four repetitions of auctions and the "Acquiring a Company Task" over four weeks with ample time for reflection about the tasks. In addition, students were given information about the bids and outcomes of other participants. In line with previous studies, none of the manipulations helped participants to avoid the curse.

More promising results regarding the ability to overcome the winner's curse were reported by Ball *et al.* (1991) and Bereby-Meyer and Grosskopf (2005). Ball *et al.* (1991) had participants reverse roles in order to make the information asymmetry more salient to the buyers. For the group that switched roles, the number of learners (defined as participants who bid zero from some trial on until the end of the experiment) increased from 9 percent to 37 percent. However, those not defined as learners only reduced their mean bid by a small amount, from 50 to 34. Moreover, changes in bidding generally occurred at the beginning of the second session in which roles were reversed, and almost no adjustment occurred during the second session. Approaching the problem from an

adaptive learning perspective, Bereby-Meyer and Grosskopf (2005) show that a decline in bids can be achieved by reducing the variability in the feedback that participants receive. However, no indication of a deep understanding of the task was observed; i.e., participants drifted in the correct direction, but did not understand the adaptation well enough to decide not to bid.

Overall, the "Acquiring a Company Task" has demonstrated surprising robustness, refuting many of the early criticisms that the observed irrational behavior would be easily corrected with experimental variation. This robustness led us to pursue an alternative approach by applying findings of the psychological literature on feedback and its implications for learning to the "Acquiring a Company Task."

Most of the research on learning in the "Acquiring a Company Task" reviewed above gave participants feedback only on the outcome of the task, i.e., the value of the company and the profit for the buyer. In contrast, Doherty and Balzer (1988) argue that feedback is the process by which an environment returns to individuals a portion of the information in their response output necessary to compare their present strategy with a representation of an ideal strategy. The psychological literature on feedback identifies two main categories: *outcome* feedback versus *cognitive* (Balzer, Doherty, O'Connor, 1989) or *process* feedback (Hammond, Stewart, Brehmer and Steinmann, 1975). Outcome feedback is simply knowledge of the outcome of a decision. In contrast, cognitive or process feedback refers to information about relations in the environment rather than outcomes.

Outcome feedback tends to be less effective in improving accuracy and performance than cognitive or process feedback (Balzer *et al.*, 1989; Hammond,

Summers and Deane, 1973); this is especially true for complex tasks (Brehmer, 1980; Hammond *et al.*, 1975; Hoffman, Earle and Slovic, 1981). Moreover, for feedback to directly improve learning, it must help the recipient to reject erroneous hypotheses (Kluger and DeNisi, 1996). In the "Acquiring a Company Task," even if a participant behaves sub-optimally, she can gain a positive amount of money approximately 33 percent of the time her offer is accepted (see Bereby-Meyer and Grosskopf, 2005). If participants are focusing on outcome feedback, it is difficult for them to make the connection between the information asymmetry and its effects on their payoffs. Consequently, the resistance of the sub-optimal behavior in the "Acquiring a Company Task" to various learning manipulations may be the result of the type of feedback given to participants.

We therefore added cognitive feedback to the "Acquiring a Company Task" as suggested by the psychology literature. The cognitive feedback was provided based on the two key parameters that drive bidding behavior in the "Acquiring a Company Task.": (1) the fact that the company is worth more to the buyer than to the seller and (2) the ignorance of the impact of the information asymmetry. In our experimental setup we therefore manipulated either (1) the additional value of the company to the buyer or (2) the information setup (asymmetry/symmetry). We hypothesized that presenting participants with variations on the task elements that were found to impede performance and having participants compare and contrast the different scenarios would trigger a deep understanding of the task and teach participants to overcome the bias.

Using contrasts or highlighting differences has been found to facilitate perceptual learning (Bransford, Franks, Vye, and Sherwood, 1989), the learning of problem-solving

principles (VanderStoep and Seifert, 1993), and concept acquisition (Tennyson, 1973). In the analogical reasoning literature it has been demonstrated that a “near-miss” contrast (Winston, 1975), or a contrasting example that differs in one important aspect from the study problem at hand, is effective in helping people abstract a solution schema and subsequently apply it to an analogous test problem (Gick and Paterson, 1992; Ross and Kilbane, 1997).

Our paper reports two experiments that show that participants continue to overbid even when given cognitive feedback in addition to outcome feedback. Experiment 1 shows that bidders usually bid higher when their premium for acquiring the company is higher. Furthermore, bidding behavior in trials when the seller and the buyer had the same amount of information (i.e., symmetric information) was similar to the bidding behavior when the seller had more information (i.e., asymmetric information), even when participants were explicitly asked to compare and contrast the two conditions.

Experiment 2 takes our manipulations a step further by adding information about foregone payoffs, i.e. payoffs of bids that were not chosen, to the cognitive feedback. This manipulation aimed to emphasize the negative expected payoff from bidding by explicitly providing information about what a participant could have received by doing something else. We additionally replaced the computer by a real seller hoping that this would increase the chance that participants took the perspective of others into account. This hypothesis was based on the growing literature on fairness that suggests that people tend to compare their payoff to others (e.g., Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000). We hypothesized that having a real seller will lead a buyer to compare his/her payoff to the seller's payoff. This, in turn, should lead to the realization that the

poor performance is driven by the fact that the seller has more information. However, in this experiment too, participants continued to overbid.

These two experiments together demonstrate the robustness of the winner's curse phenomenon in the "Acquiring a Company Task" in the face of cognitive learning manipulations. The paper is organized as follows: Section 2 presents the experimental design and results of Experiment 1. Section 3 does the same for Experiment 2. Section 4 concludes.

2. Experiment 1

Participants: Ninety participants recruited from the Boston area, aged 18 to 60, participated in this study and were paid according to their performance. They received a $10 show-up fee and had the opportunity to earn up to an additional $10.

Procedure: The computerized experiment was run in the experimental laboratory at Harvard Business School and included four treatments: (1) Control, (2) Varying *k*, (3) Asymmetric-Symmetric, and (4) Asymmetric-Symmetric with Comparison.[1] In all treatments of Experiment 1, participants played the role of buyers in the "Acquiring a Company Task" for 100 trials, divided into 5 parts of 20 trials each. In each trial, in each of the parts, participants were asked to determine a price they would like to offer for the company. They were told that the value of the company is distributed uniformly between 0 and 100, and that a computer program represents the seller. At the beginning of the

[1] Conditions 1-3 were run simultaneously and participants were assigned randomly to the different conditions. Condition 4 was run later and aimed to strengthen the manipulation of condition 3. The same subject pool and recruiting procedures were used. Note that payments differ slightly. While participants in

experiment each participant received an endowment of 1,500 points worth 0.7 cents each. The feedback after each trial included (1) the realized value of the company, (2) the value of the company to the buyer, and (3) the buyer's profit/loss.[2]

Treatment 1: Control

In the Control treatment the company was always worth 50 percent more under the buyer's management than under the seller's management. We refer to this additional value as the *k*-value. The seller knew the true value of the company before bids were evaluated. The buyer did not know the company's true value; hence the information setup was asymmetric.

Treatment 2: Varying *k*

In the Varying *k* treatment, the basic procedure was the same as in the Control treatment, except for the fact that the values of *k* alternated from one part of 20 trials to the other. Half of the participants in this treatment saw a sequence of: 1.2 - 2.5 - 1.2 - 2.5 followed by 1.5, and the other half saw the reversed order: 2.5 - 1.2 - 2.5 - 1.2 followed by 1.5. The last part of 20 trials was kept the same over all treatments (k = 1.5, asymmetric information) in order to evaluate performance and compare the different treatments. Participants were informed that the value of *k* would vary in each part and that they would be informed about the exact value prior to starting a new part. In the Varying *k* treatment the setting was asymmetric just like in the Control treatment. Clearly, when k = 1.2, participants should not bid, and when k = 2.5, the payoff maximizing choice is to

conditions 1-3 were paid for all 100 rounds at an exchange rate of 0.7 cents per point, participants in condition 4 were only paid in the last 60 rounds at an exchange rate of 1.2 cent per point.

[2] A sample set of instructions can be found in the appendix.

bid 100.[3] The idea of this treatment was to help participants see the importance of this parameter when they made the final 20 decisions.

Treatment 3: Asymmetric-Symmetric

In the Asymmetric-Symmetric treatment the procedure was the same as in the Varying *k* treatment, except that instead of varying the additional value to the buyer (*k*), we varied the information setting. Participants were informed about the fact that sometimes the seller will know the exact value of the company before bids are evaluated and sometimes he will not. Here, too, we ran two orders: one that started with 20 trials of the asymmetric setting and one that started with 20 trials of the symmetric setting. In both information settings, however, the value of the company to the buyer was kept constant at $k = 1.5$. And, as before, both orders ended with 20 asymmetric trials, to be compared to the Control treatment.

Treatment 4: Asymmetric-Symmetric with Comparison

The basic procedure in this treatment was similar to the Asymmetric-Symmetric treatment, except for the following changes. First, we gave participants only one sequence, Asymmetric – Symmetric – Asymmetric – Symmetric – Asymmetric, since we did not find any difference resulting from the order that participants experienced the different conditions. After every 40 trials (after the completion of two different parts of 20 trials each), participants saw the *full history* of their offers, the value of the companies, and their profit in each trial. The computer screen was divided: on one side participants

[3] When the value of the company to the buyer is 20 percent more than the value of the company to the seller, the expected value of any accepted positive bid is negative, as demonstrated by the following equation: $E(bid/bid>v) = 1.2*E(v|bid>v) – bid = 1.2*bid/2 – bid = -0.4*bid$. Consequently, the optimal bid for a risk-neutral bidder is zero. By contrast, when the company is worth 150 percent more to the buyer

received feedback for the symmetric information trials and on the other side they received feedback for 20 trials of the asymmetric information; we did this to make the differences between the two conditions more salient. To further help participants compare the different conditions, we asked two questions after every 40 trials: (1) "What is the key difference between the two parts?" and (2) "Identify an overall principle that should guide your decision regarding how much to offer in each part." Based on the aforementioned evidence from the psychology literature we hypothesized that these questions would facilitate the extraction of the general rule, "do not bid in the asymmetric information setting with $k = 1.5$." The first 40 trials were unpaid learning trials in order to give participants the opportunity to explore the game without risking money and to weaken the focus on outcome feedback. In the remaining 60 trials, participants received 1,500 points as endowment and each point was worth 0.012 cents.

2.1. Experimental Results

An Anova analysis on the mean bid in the last part (last 20 trials of 100) as a dependent variable and the treatment as an independent variable revealed no significant difference between treatments ($F(3, 86)=1.56$, $p<0.2$, $MSe=467.6$). The mean and standard deviations of the bids in the last 20 trials ($k=1.5$, asymmetric information) are presented in Table I.

To deepen our understanding of the learning process, we examined the mean bids per block of 20 trials in the different treatments. Figure 1 presents the mean bids and

than to the seller, the expected value of any accepted positive bid is positive, i.e., $E(bid/bid>v)= 2.5*E(v|bid>v) – bid = 2.5*bid/2 – bid = 0.25*bid$. Consequently, in this case it is optimal to bid 100.

their standard deviation in the five parts of the Control treatment.[4] As can be seen, almost no adaptation (towards choices of zero) occurred as a function of experience. Neither the mean nor the variance changed as a function of experience.

Treatment	Mean	Standard Deviation
Control (N = 20)	38.84	23.91
Varying *K* (N = 20)	45.41	20.89
Asymmetric-Symmetric (N = 15)	52.84	17.59
Asymmetric-Symmetric with Comparison (N = 35)	40.14	22.18

Table I: Mean and standard deviations of bids in the last 20 trials in Experiment 1
(k=1.5, asymmetric information, number of observations, N, is given in parentheses)

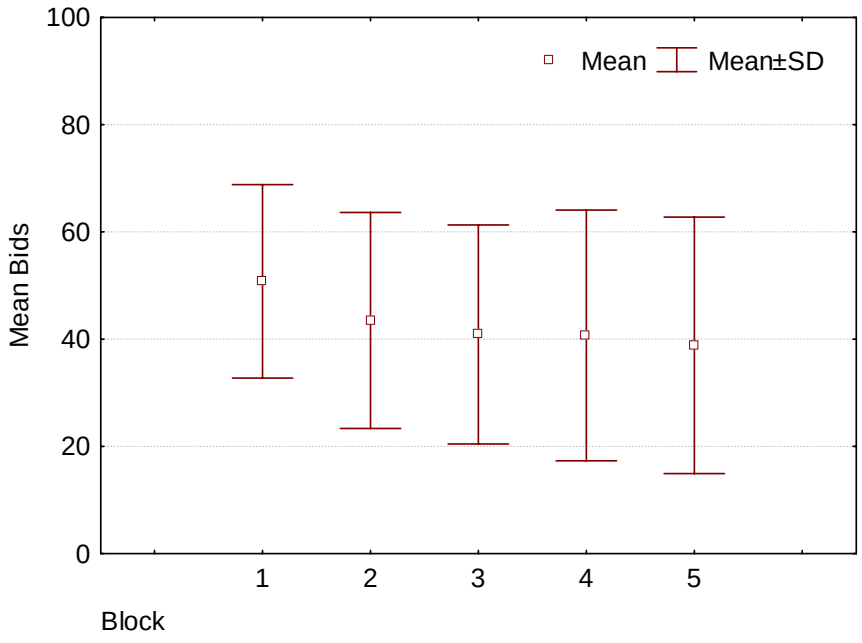

Figure 1: Mean bids and their standard deviation in the Control treatment
(blocks of 20 trials each)

[4] The variance refers to the variation of the individual means. The individual means are calculated for blocks of 20 trials each.

Figure 2 presents the mean bids per block of 20 trials for the Varying *k* treatment for both sequences. As can be seen, participants differentiated between the different *k*–values; their bids were higher when $k = 2.5$ than when $k = 1.2$. Yet the bids were too low when $k = 2.5$ (less than 100) and too high (greater than zero) when $k = 1.2$ when compared to the rational prediction. When participants faced a *k*–value of 1.5 in the last 20 trials, their mean bids were higher than were the mean bids when $k = 1.2$. Observing that bids are a function of the *k*–value leads us to believe that participants did not gain a deep understanding of the general principle. A deeper understanding of the task would manifest itself in no difference in mean bids across a *k*–value of 1.2 and 1.5, since the rational answer is zero in both cases. No difference was found in the mean bids in the last 20 trials of the different sequences.

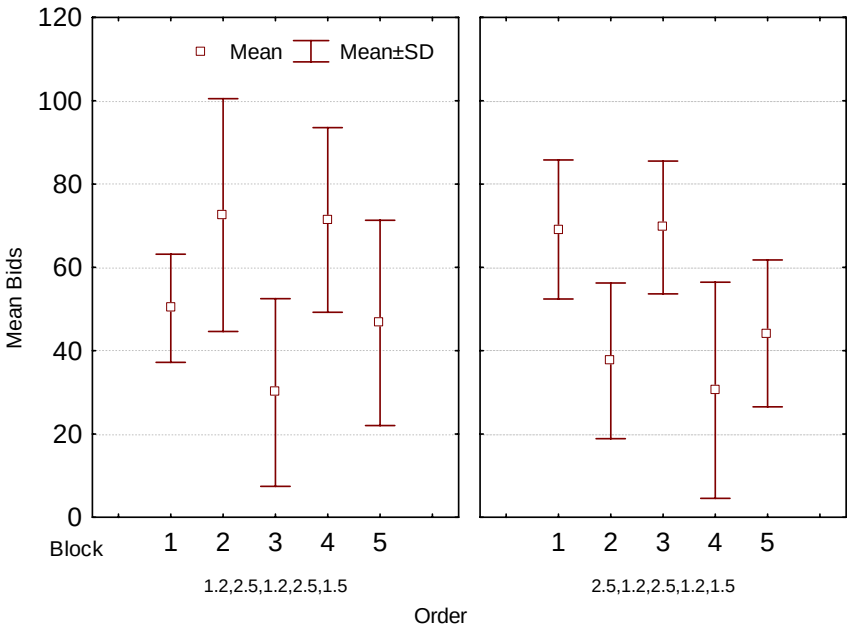

Figure 2: Mean bids in the Varying *k* treatment
(blocks of 20 trials each)

Note that the observed underbidding when k = 2.5 is consistent with risk averse preferences. However, observed overbidding when k = 1.2 would indicate risk loving preferences. Since these are within subject variations, we do not think that risk preferences can explain our data.

Figure 3 presents the mean bids in each block of 20 trials in the Asymmetric-Symmetric treatment for the two sequences. As can be seen, there is almost no difference between the mean bids in the different parts. Participants did not differentiate between the symmetric and the asymmetric setting. The mean bids in the last 20 trials were significantly higher than zero and not significantly different from the mean bids in the Control treatment. These results show that experiencing different seller information settings does not help participants to avoid the winner's curse. This lack of difference is broadly compatible with prior conclusions that participants in the asymmetric Acquiring a Company task treat it as knowledge were symmetric (Samuelson and Bazerman, 1985).

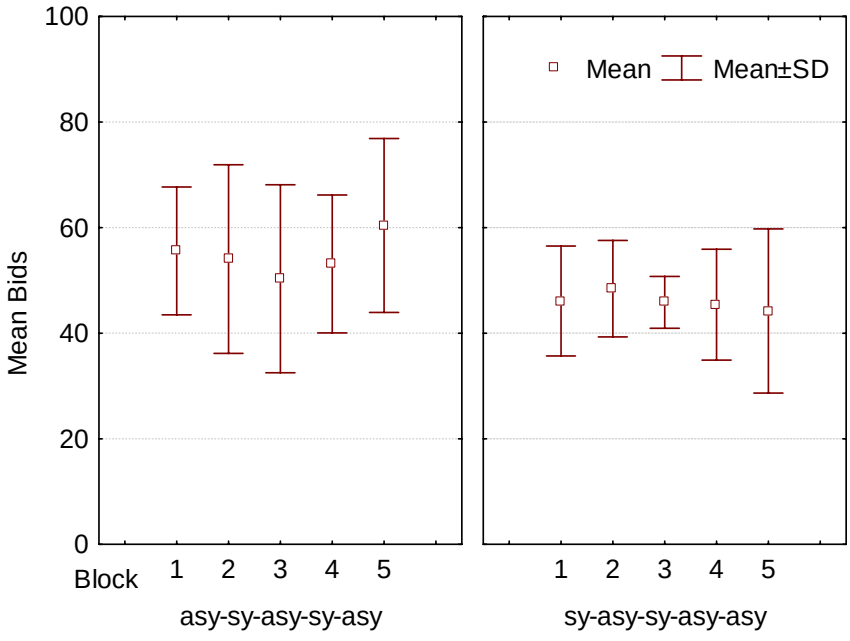

Figure 3: Mean bids in the Asymmetric-Symmetric treatment
(blocks of 20 trials each)

Figure 4 shows mean bids and their standard deviation in each block of 20 trials in the Asymmetric-Symmetric with Comparison treatment. Here too, subjects did not seem to differentiate between the two information settings even after being asked to compare and contrast the two and after receiving full information regarding their bid history.

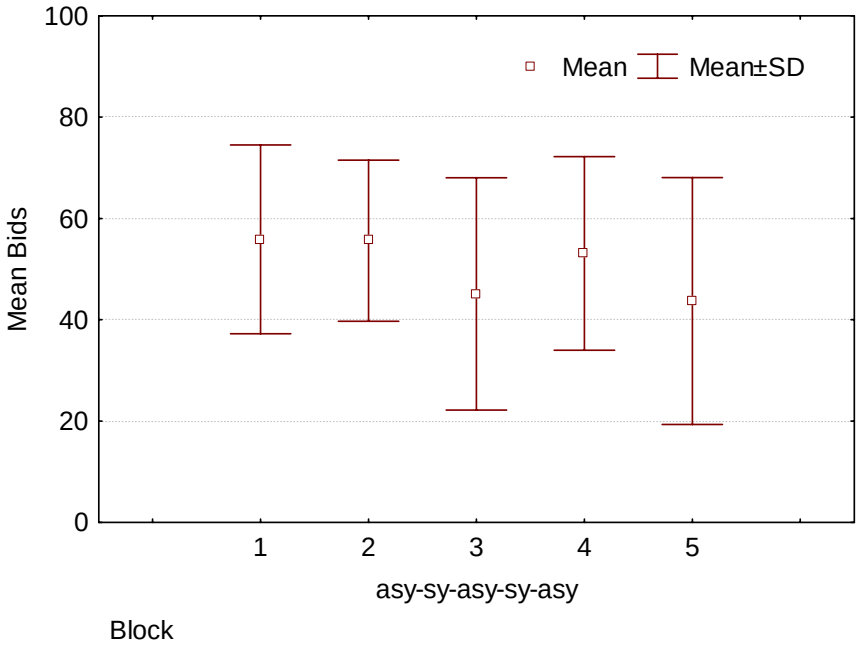

Figure 4: Mean bids and standard deviations
in the Asymmetric-Symmetric with Comparison treatment
(blocks of 20 trials each)

In summary, our results reveal that neither changing the *k*-value nor changing the information setting helps participants overcome the winner's curse. Although participants bid differently depending on the value of *k*, they do not pay attention to the information setting and specifically to whether the seller has more information than they have themselves. As a result, bidding behavior in the symmetric case was similar to bidding behavior in the asymmetric case. The inability to differentiate between the two settings persisted even when participants were explicitly asked to compare the two.

3. Experiment 2

To make the negative expected profit from bidding in the Asymmetric-Symmetric treatment more salient, we gave participants feedback on the payoff they would have gotten if they had submitted any other bid; i.e., we provided foregone payoff information given the realization of the value of the company. Note that while participants in Experiment 1 could compute this information, in Experiment 2 it was explicitly provided. This foregone payoff feedback should help participants realize that even if they won by submitting a positive bid in a specific trial, on average, it is not worthwhile to bid. While economists have started to look at the effects of foregone payoff information on learning (e.g. Camerer and Ho, 1999; Grosskopf, Erev and Yechiam, 2004; Heller and Sarin, 2001), no clear pattern has evolved. We are introducing foregone payoff information as a cognitive feedback manipulation. In addition, we had real sellers for the first 80 trials. Those real sellers were then assigned the role of buyers for the last 20 trials (against a computer).[5] We were interested in two things: (1) does buyer behavior change if the money they are bidding (losing) goes to another participant instead of staying with the experimenter, and (2) does playing the role of the seller before playing the role of the buyer make the asymmetric information more obvious to the buyer?[6]

We also wanted to have a clear-cut comparison of what participants learned in the experiment. Therefore, before participants had any experience with the task (and before they even knew they would participate in 100 trials), we asked them what they thought

[5] Additionally, the last 20 trials are used for comparison to the other treatments.
[6] There is a caveat though. Adding a human participant on the seller's side adds complexity since buyers had to think about what sellers would be doing instead of being told that a computer program would accept or reject bids according to a specified rule.

the optimal bid was. At the end of the experiment, after 100 trials, we asked participants what they would advise a friend to bid in such a setting. The ability to transfer learned skills to other people has been suggested as one of the main factors that distinguishes experience from expertise (Bazerman and Neale, 1993). Mere experience is often not enough to critically analyze the learned situation. To transform experience into expertise, the learner must understand what she experienced, distinguish what makes successful experiences different from unsuccessful experiences, and develop a conceptual understanding of the process. This understanding in turn will enable the transfer of knowledge to other people. Consequently, studying the ability to give advice will enable us to gain knowledge about the level of expertise acquired by participants.[7] In addition, giving advice allows participants to reveal their knowledge of the best strategy, a knowledge they might not have been able to reveal during the course of the experiment due to risk-taking behavior they could have developed in the loss domain (Kahneman and Tversky, 1979).

Participants: Twenty students of industrial engineering at Ben Gurion University in Israel participated in the experiment for course credit. In addition, they could earn an average of 10 Israeli Shekels (around $2.50) depending on their performance.

Procedure: The experiment included two treatments: "Control" and "Foregone payoff." The basic procedure of the two treatments was similar to the Asymmetric-

[7] In contrast to studies on how advice affects behavior in intergenerational settings (e.g. Schotter and Sopher, 2001), our advice is purely hypothetical and does not have any payoff relevance. Moreover, we focus on the adviser's behavior, whereas studies of intergenerational settings mostly explore advisee's behavior.

Symmetric with Comparison treatment in Experiment 1 (treatment 4), except for the fact that the seller was a human participant and not a computer program.

Participants were randomly assigned to one of the two treatments, and were assigned to be either sellers or buyers for 80 trials (five sellers interacted with five buyers in each treatment). They were re-matched after every trial with a different partner. In the last 20 trials all participants played the role of the buyer against a programmed seller, as was the case for the four treatments in Experiment 1. Participants' monetary payoff was determined by their performance in one random trial out of the last twenty trials. The exchange rate was 1 Shekel ($0.20) for 10 points.[8]

In the "Foregone Payoff" treatment, after each trial, a table on the right hand side of the computer screen told participants what points they would have received if they had made a different bid, assuming they were matched with a profit-maximizing seller. For example, if a participant submitted a bid of 80, and the realized value of the company was 40 (resulting in a loss of 20), the participant would have seen what payoff (profit or loss) she would have received if she had submitted a different bid. Bid values of 0 to 100 were shown, holding the realized value of the company fixed.

In both treatments, at the beginning of the experiment, after reading the "Acquiring a Company Task" but before experiencing it or finding out that they would engage in 100 trials of the task, participants were asked to supply the optimal bid. Similarly, at the end of the experiment, participants were asked to advise a friend who

[8] Note that in Experiment 1 participants were paid for either all rounds (treatments 1-3) of for the last 60 rounds (treatment 4). Paying one round at random eliminates potential income effects. As for the overall smaller payment in Israel, we followed what is usually paid in experiments and do not believe that the smaller payments influence behavior differently in comparison to higher payments in the United States.

was about to participate in the same experiment under the same conditions on the optimal offer. They were asked to give advice under two conditions: 1) the seller knows the value of the company and 2) the seller does not know the value of the company. Participants were not paid neither for their evaluation of the optimal bid nor for their advice.

3.1. Results

Table II shows mean optimal bid evaluation before Experiment 2 and the advice that was given after Experiment 2.

	Evaluation *before* Experiment	Advice *after* Experiment	
Control	79.22 (SD=30.31)	Asymmetric	47.5 (SD=25.05)
		Symmetric	49.85 (SD=9.90)
Foregone	72.60 (SD=31.39)	Asymmetric	54.38 (SD=24.99)
		Symmetric	58.33 (SD=13.46)

Table II: Mean Evaluation/Advice in Experiment 2
(standard deviations are given in parentheses)

The evaluation before the experiment, as expected, does not differ significantly across treatments ($t(17)$=-0.55, p=0.13.). For both treatments, the advice given after the experiment is smaller than their initial evaluation, ($t(15)$=2.93, p=0.01 (Control-Asym), $t(15)$=2.46, p=0.02 (Forgone-Asym). While participants in the Control and in the Foregone treatment lower their bids in response to having gained experience, it remains questionable whether they understood the underlying principle. In particular, the advice

given does not differ between the asymmetric condition and the symmetric condition in neither of the treatments. This shows that participants did not acquire a deep understanding of the task even though their overall advice is smaller than their initial evaluation.

Figure 5 presents the mean bids for every block of 20 trials. For the first 80 trials, the results present the mean bids of the buyers only. In the last 20 trials, since all participants played the role of buyers, the figure presents the results of participants who previously were in the role of buyers and sellers separately.

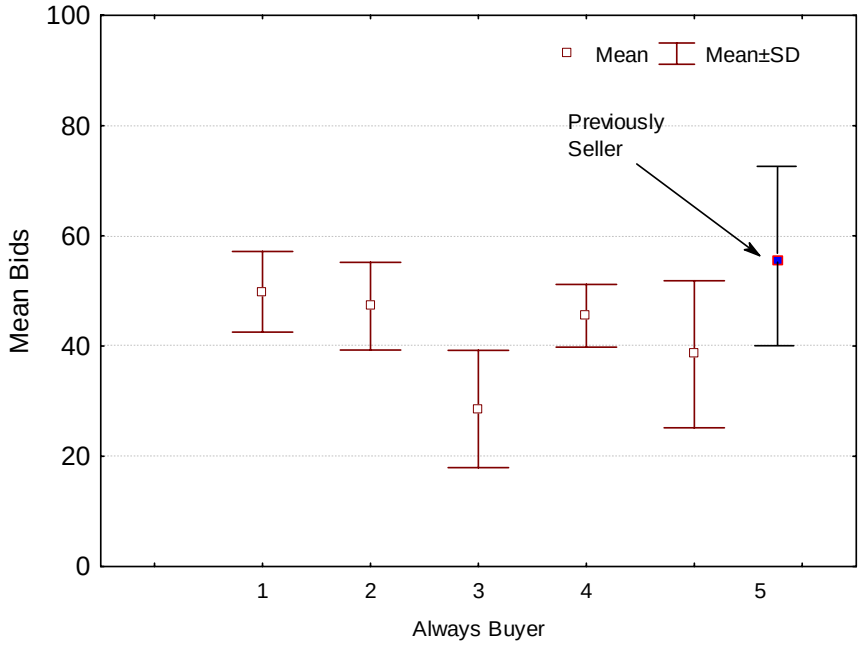

Figure 5: Mean bids in the Control Treatment of Experiment 2
(last block of 20 trials separates participants who were previously buyers/sellers)

As can be seen, no improvement in learning was found in this treatment compared to the results of the "Asymmetric-Symmetric with Comparison" treatment in Experiment 1 with a programmed opponent. The mean bid (of participants who always were buyers) in the last 20 trials was not significantly lower than the mean bid at the beginning.

Moreover, the experience as sellers did not help participants in the part in which they became buyers. On the contrary, their performance was marginally worse than the performance of the buyers who kept their roles throughout the entire duration of the experiment ($t(8)=-1.84$, $p=0.1$).

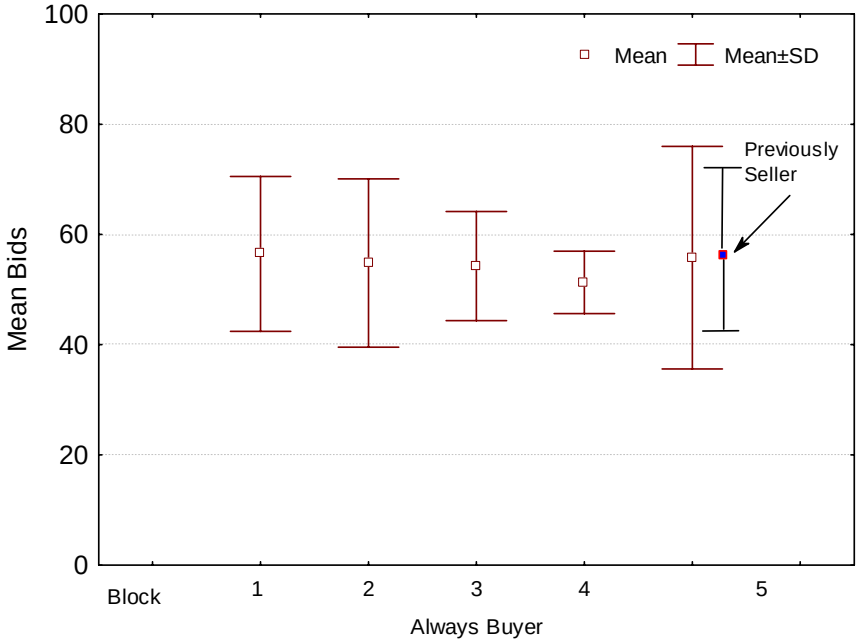

Figure 6: Mean bids in the Foregone treatment of Experiment 2
(last block of 20 trials separates participants who were previously buyers/sellers)

Figure 6 presents the mean bid in the Foregone treatment for every block of 20 trials. As can be seen, no improvement in learning was found in this treatment compared to the Control treatment when the foregone payoff information was not explicitly provided. Additionally, we observe that the behavior of participants who had different roles (buyers or sellers respectively) for the first 80 trials no longer differed.

The current results show that even when participants are given information regarding the expected payoff for every other possible bid, i.e. foregone payoff information, participants remain unable to learn not to bid.

4. Discussion

A common bias in decision-making is the inability of decision makers to adjust their behavior in anticipation of their opponents' choices and the resulting consequences. This bias can be very costly, as demonstrated when winners of auctions or bilateral negotiations receive a prize whose cost exceeds its value, a phenomenon known as the winner's curse. Researchers have used the "Acquiring a Company Task" extensively to demonstrate the winner's curse and to analyze barriers to overcoming the underlying bias. An important question is whether it is possible to learn to overcome the winner's curse. A number of studies now point to the robustness of the phenomenon.

Research on the effect of learning in this task reveals some adaptation of behavior towards the optimal solution with no indication of a deep understanding of the problem. In most of the studies that examined learning in the "Acquiring a Company Task," participants received mere outcome feedback, which has been found to be ineffective at improving performance in complex tasks. In the current study, in addition to outcome feedback we gave participants process (i.e. cognitive) feedback that aimed to deepen their understanding of the task and has been shown to improve performance in complex tasks. By isolating and separately manipulating the two main barriers to improved decision making in this task, we let participants experience a situation in which it is not optimal to bid, followed by one in which it is optimal to submit positive bids. Based on the problem-solving literature, we expected that by presenting bidders with these two situations they would extract the general principles for optimal bidding. However, we did not find any improvement in participants' performance, not even when they were explicitly asked to compare and contrast the different scenarios (Experiment 1). Giving

participants information about their foregone payoff additionally to the outcome payoff and having a real seller who cashes in the lost money of the bidder does not improve performance either (Experiment 2).

Not so long ago, experimental researchers believed that the process of improving judgment would occur naturally as individuals receive feedback about their past decisions. For example, Kagel and Levin (1986, p. 917) wrote:

> *"Given sufficient experience and feedback regarding the outcomes of their decisions, we have no doubt that our experimental participants, as well as most bidders in "real world" settings, would eventually learn to avoid the winner's curse in any particular set of circumstances. The winner's curse is a disequilibrium phenomenon that will correct itself given sufficient time and the right kind of information feedback."*

It is clear that experiences help us to improve numerous skills. However, Kagel and Levin's assumption may not be so obvious. At about the same time, Tversky and Kahneman (1986, pp. 274 – 275) argued that

> *"Basic judgmental biases are unlikely to correct themselves over time, since responsive learning requires accurate and immediate feedback, which is rarely available in the real world because: (i) outcomes are commonly delayed and not easily attributable to a particular action; (ii) variability in the environment degrades the reliability of feedback. . . .; (iii) there is often no information about what the outcome would have been*

> *if another decision had been taken; and (iv) most important decisions are unique and therefore provide little opportunity for learning (see Einhorn and Hogarth, 1978) . . . any claim that a particular error will be eliminated by experience must be supported by demonstrating that the conditions for effective learning are satisfied.*"

The data in the current experiment point to the predictive ability of Tversky and Kahneman over Kagel and Levin, at least as far the "Acquiring a Company Task" is concerned. We conclude that the winner's curse observed in the "Acquiring a Company Task" is very robust. It seems as if the fact that sub-optimal choices (bids greater than zero) sometimes lead to positive payoffs imposes yet another barrier to overcoming the underlying bias. The tendency of decision-makers to adaptively make their choices based on past experience impedes the effectiveness of cognitive feedback.

While the value of learning is too important to abandon, we must conclude that the "Acquiring a Company Task", which requires little or no mathematical sophistication, is unlikely to be solved using conventional manipulations of experience. Building off the work of Thompson *et al.*, (2000) on analogical reasoning, some current research suggests that having participants compare and contrast their choices within and across problems might shed light on effective learning mechanisms to improve decision making in the "Acquiring a Company" task and similar problems (e.g. Idson *et al.*, 2003).

References

Akerlof, G. (1970). The market for lemons: Quality uncertainty and the market mechanism, *Quarterly Journal of Economics*, **84**, 488 – 500.

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.

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

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

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

Bazerman, M. H. (2002). *Judgment in Managerial Decision Making* (5th Ed.). New York: John Wiley & Sons.

Bazerman, M. H., and W.F. Samuelson (1983). I won the auction but don't want the prize. *Journal of Conflict Resolution*, **27**, 618-34.

Bazerman, M. H. and M.A. Neale (1992). *Negotiating Rationally*. New York: The Free Press.

Bereby-Meyer, Y., and B. Grosskopf (2005). Overcoming the Winner's Curse: An Adaptive Learning Perspective, AoM Conflict Management Division 2002 Mtgs. No. 13496. http://ssrn.com/abstract=324201.

Bransford, J.D., Franks, J.J., Vye, N.J., and R.D. Sherwood (1989). New approaches to instruction: Because wisdom can't be told. In S. Vosniadou and A. Ortony (Eds.),

Similarity and analogical reasoning (pp. 470-497). Cambridge, England: Cambridge University Press.

Bolton, G. and A. Ockenfels (2000), "ERC: A Theory of Equity, Reciprocity and Competition," *American Economic Review*, **90**, 166-193.

Brehmer, B. (1980). In one word: Not from experience. *Acta Psychologica*, **45**, 223-241.

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

Camerer, C. F., and T.H. Ho (1999). "Experience-Weighted Attraction (EWA) Learning in Normal-Form Games." *Econometrica*, 67, 827-74.

Coursey, D. L., Hovis, J. L., and W.D. Schulze (1987). The disparity between willingness to accept and willingness to pay measures of values. *The Quarterly Journal of Economics*, **102**, 679-690.

Doherty, M.E. and W.K. Balzer (1988). Cognitive feedback. In B. Brehmer and C. R. B. Joyce (Eds.), *Human Judgment: The SJT approach*. Amsterdam: North-Holland.

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

Fehr, E. and K. Schmidt, (1999), "A Theory of Fairness, Competition and Cooperation," *Quarterly Journal of Economics*, **114**, 817-868.

Fischbacher, U. (1999), z-Tree - Zurich Toolbox for Readymade Economic Experiments - Experimenter's Manual, mimeo

Foreman, P. and J.K. Murnighan (1996). Learning to Avoid the Winner's Curse. *Organizational Behavior and Human Decision Processes*. **67**, 170-180.

Gick, M.L., and K.J. Paterson (1992). Do contrasting facilitate schema acquisition and analogical transfer? *Canadian Journal of Psychology*, **46**, 539-550.

Grosskopf, B., I. Erev and E. Yechiam (2006), "Foregone with the Wind," *International Journal of Game Theory*, forthcoming.

Grosskopf, B. and R. Nagel (2001). Rational Reasoning or Adaptive Behavior? Evidence from Two-Person Beauty – Contest - Games, Harvard NOM Research Paper No. 01-09. http://ssrn.com/abstract=286573.

Hammond, K. R., Stewart, T. R., Brehmer, B., and D.O. Steinmann (1975). Social judgment theory. In Kaplan and S. Schwartz (Eds.), *Human Judgment and Decision Processes*. New York: Academic Press.

Hammond, K. R., Summers, D. A., and D.H. Deane (1973). Negative effects of outcome-feedback in multiple-cue probability learning. *Organizational Behavior and Human Performance*, **9**, 30-34.

Heller, D.and R. Sarin (2001), "Adaptive Learning with Indirect Payoff Information," mimeo.

Hoffman, P. J., Earle, T. C., and P. Slovic (1981). Multidimensional functional learning (MFL) and some new conceptions of feedback. *Organizational Behavior and Human Performance*, **27**, 75-102.

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

Kagel, J. H. (1995). Auctions: A survey of experimental research. In Kagel, J. H. and Roth, A. E. (Eds). *The Handbook of Experimental Economics*.

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

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

Kluger, A. N. and A. DeNisi (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.

Ross, B.H., and M.C. Kilbane (1997). Effects of principle explanation and superficial similarity on analogical mapping in problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **23**, 427-440.

Roth, A. E., and I. Erev (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., and M.H. Bazerman (1985), Negotiation under the winner's curse, in V. Smith (ed.) *Research in Experimental Economics*, Vol. III, Greenwich, CT, JAI Press.

Schotter, A. and B. Sopher (2001), Advice and Behavior in Intergenerational Ultimatum Games: An Experimental Approach, mimeo.

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

Smith, V. L. (1985). Experimental Economics: Reply. *American Economic Review,* **75**, 265-272.

Tennyson, R.D. (1973). Effect of negative instances in concept acquisition using a verbal learning task. *Journal of Educational Psychology*, **64**, 247-260.

Thompson, L., Gentner, D., and J. Loewenstein (2000). Analogical training more powerful than individual case training. *Organizational Behavior and Human Decision Processes*, **82**, 60-75.

Tor, A., and M.H. Bazerman (2003). Focusing Failures in Competitive Environments: Explaining Decision Errors in the Monty Hall Game, the Acquiring a Company Problem, and Multi-party Ultimatums, *Journal of Behavioral Decision Making*, **16**, 353-374.

Tversky, A. and D. Kahneman (1974). Judgment under uncertainty: Heuristics and biases. *Science*, **185**, 251-284.

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

VanderStoep, S.W., and C.M. Seifert (1993). Learning "how" versus learning "when": Improving transfer of problem-solving principles. *The Journal of Learning Sciences*, **3**, 93-111.

Winston, P.H. (1975). Learning structural descriptions from examples. In P.H. Winston (Ed.), *The psychology of computer vision* (pp. 157-210). New York: McGraw-Hill.

Appendix

Sample Set of Instructions (Asymmetric – Symmetric Treatment)

Welcome to the Computer Lab for Experimental Research at the Harvard Business School. Thank you for agreeing to participate in this experimental investigation on economic behavior.

Please do not touch your computer until everyone is seated, and we have read the instructions together.

Please enter your identification number (the number you received at the lab entrance) here: _____

Please read along with me as I read the instructions. There will be opportunities to ask questions if the instructions are not clear.

The experiment consists of 100 rounds divided into 5 parts of 20 rounds each. Right now, you are given a general set of instructions. The necessary specifications for each part will be given to you prior to each new part on your computer screen.

In each round you will face the following general decision problem:

You will represent Company A (the potential acquirer), which is currently considering acquiring Company T (the target). You plan to pay in cash for 100% of Company T's shares but are unsure how high a price to offer. The main complication is this: the value of Company T depends directly on the outcome of a major oil exploration project it 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/share). But if the project succeeds, the value of the company under current management could be as high as 100 points/share. All share values between 0 points and 100 points are considered equally likely. By all estimates, the company will

be worth considerately more in the hands of Company A (you) than under current management. In fact, the company will be worth 50% more under the management of A (you) than under Company T. If the project fails, the company is 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 (you) is 75 points/share. Similarly, a 100 points/share value under Company T implies a 150 points/share value under Company A, and so on.

The board of directors of Company A has asked you to determine the price they should offer for Company T's shares. This offer must be made now, before the outcome of the drilling project is known. From all indications, Company T would be happy to be acquired by Company A, provided it is at a profitable price. Moreover, Company T wishes to avoid, at all cost, the potential of a takeover bid by any other firm.

The different parts of today's experiment will vary to the extent that sometimes company T will know the results of the drilling project and sometimes it will not. You will be told at the beginning of each new part whether or not this is the case. You will be given this information on your computer screen. However, you will never know the results of the drilling project (the value of the company to its current owner) when deciding about your bid.

A computer program represents Company T. If Company T knows the value of its company, it will accept any offer that is greater than the (per share) value of the company under current management, and it will reject any bids below or equal to its true value. Thus, if you offer 60 points/share, for example, Company T will accept if the value of the company to Company T is anything less than 60 points.

If Company T on the other hand doesn't know the exact value of its company, it will accept any bids higher than its expected value of the company and reject any bids equal to or less than its expected value. The expected value of the company is calculated assuming that all values from 0 to 100 are equally likely.

As the representative of Company A, you are deliberating over price offers in the range of 0 points/share to 150 points/share. Your bid should be expressed in integer values.

Company A's assets (your assets) are currently worth 1,500 points. You will have altogether 100 trials to bid on company T in order to increase the amount of company A's (and thus your) assets. If in any trial you bid less than the value of the company under the current management (or the expected value) your bid will be rejected and your assets will neither increase nor decrease. If you bid more than the value of the company under current management (or the expected value) your bid will be accepted and your assets will change by the difference between the value under your management and your bid. Your assets can increase or decrease.

You are being paid on performance based on a rate of 7 cents for every 10 points. This means that, for example, if you end up with 1,000 points you will be paid $7.00 in addition to your show up fee. If you end up with 2,000 points you will be paid $14.00 in addition to your show up fee. Under no circumstances will you lose your show up fee.

Notice that acquiring the company is a neutral event - your performance will be judged only on the value of your assets at the end of the experiment.

There are two things, which I would like to stress before you begin the experiment. First, the instructions tell you that "you will have to bid on a company similar to company T every trial for the next 100 trials". By this we mean 100 chances to bid for different and independent companies. The true value of each company is equally likely to be any value between 0 and 100. The values of the 100 companies are randomly determined by the computer and are statistically independent. This means that knowing the value of the company or companies in any previous trial or trials gives you no information about what the value will be in the next trial. This means, for example, that if you observe a high value one trial you should neither believe that there is a trend for companies to have high values nor that the law of averages will cause the value in the next trial to be low.

Second, the instructions tell you that "acquiring a company is a neutral event - your performance will be judged only on the value of your assets at the end of the experiment". This means that you should not care about the number of companies that you acquire for any reason other than the fact that acquiring a company will either increase or decrease the value of your assets.

At the beginning of each part, you will be informed about the fact whether Company T knows the exact value of the company or not. In each trial you will have to enter the price you would like to offer into the corresponding box on the computer screen (shown below). After each trial you will be informed about the value of the company under current management, the value to you and your earnings in points.

This is what your input screen will look like:

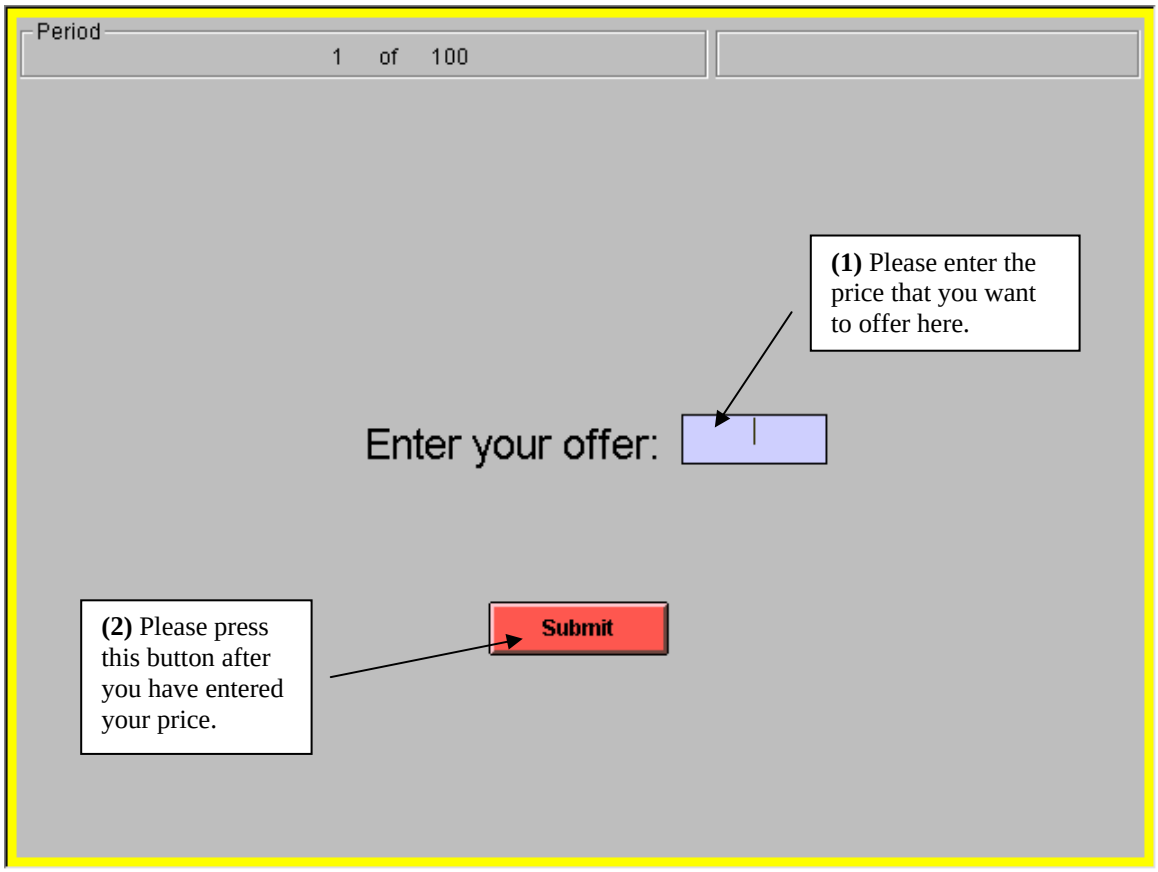

After each decision you will see the following feedback screen:

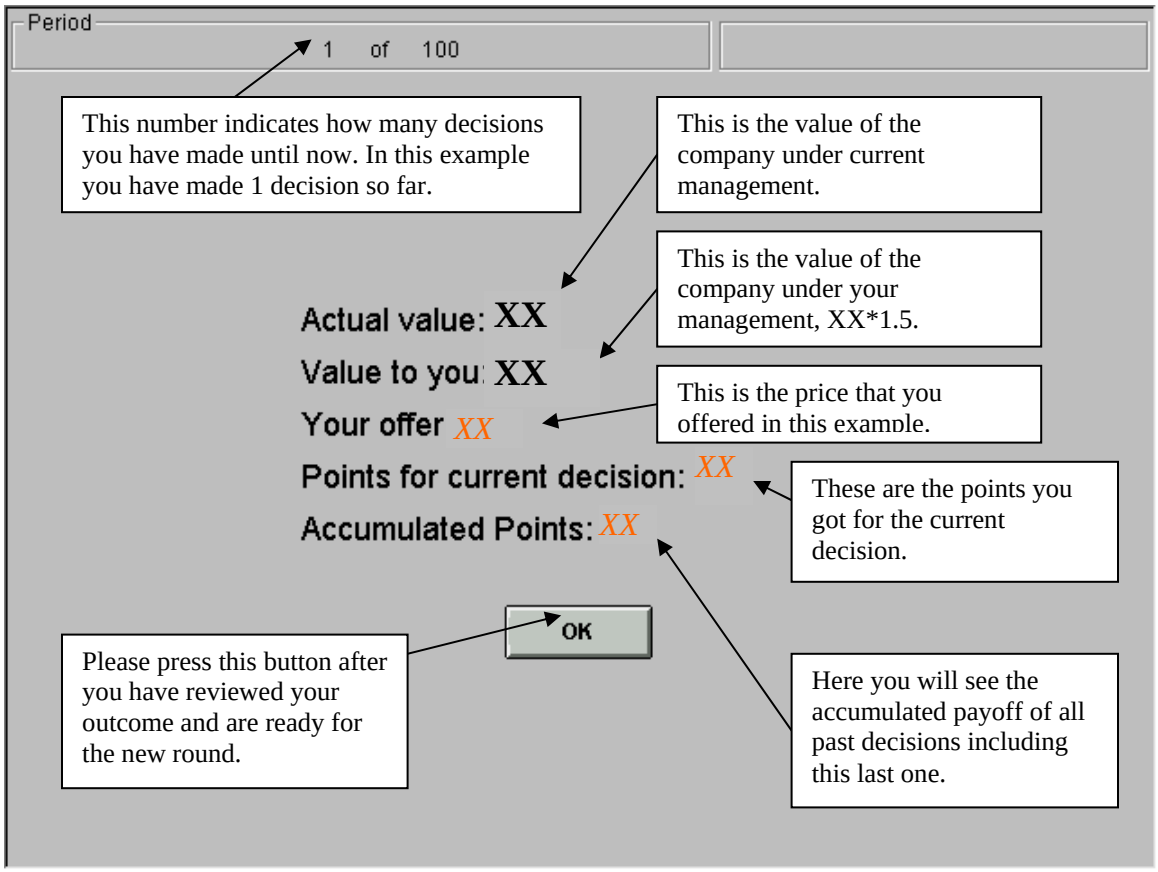

At the end of the first part (after 20 rounds) you will see the following summary screen.

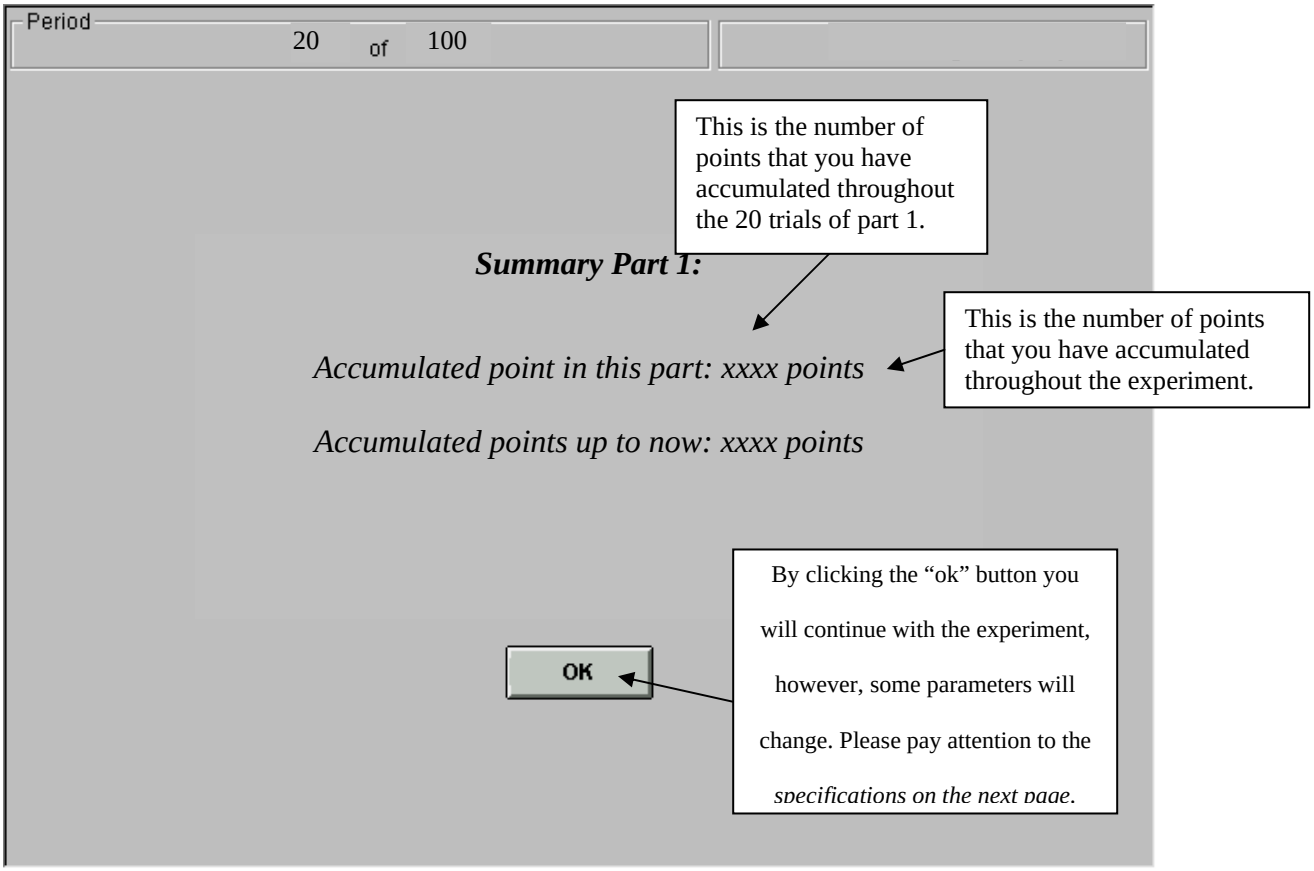

At the very end of the experiment (after 100 trials) you will be shown the final summary screen.

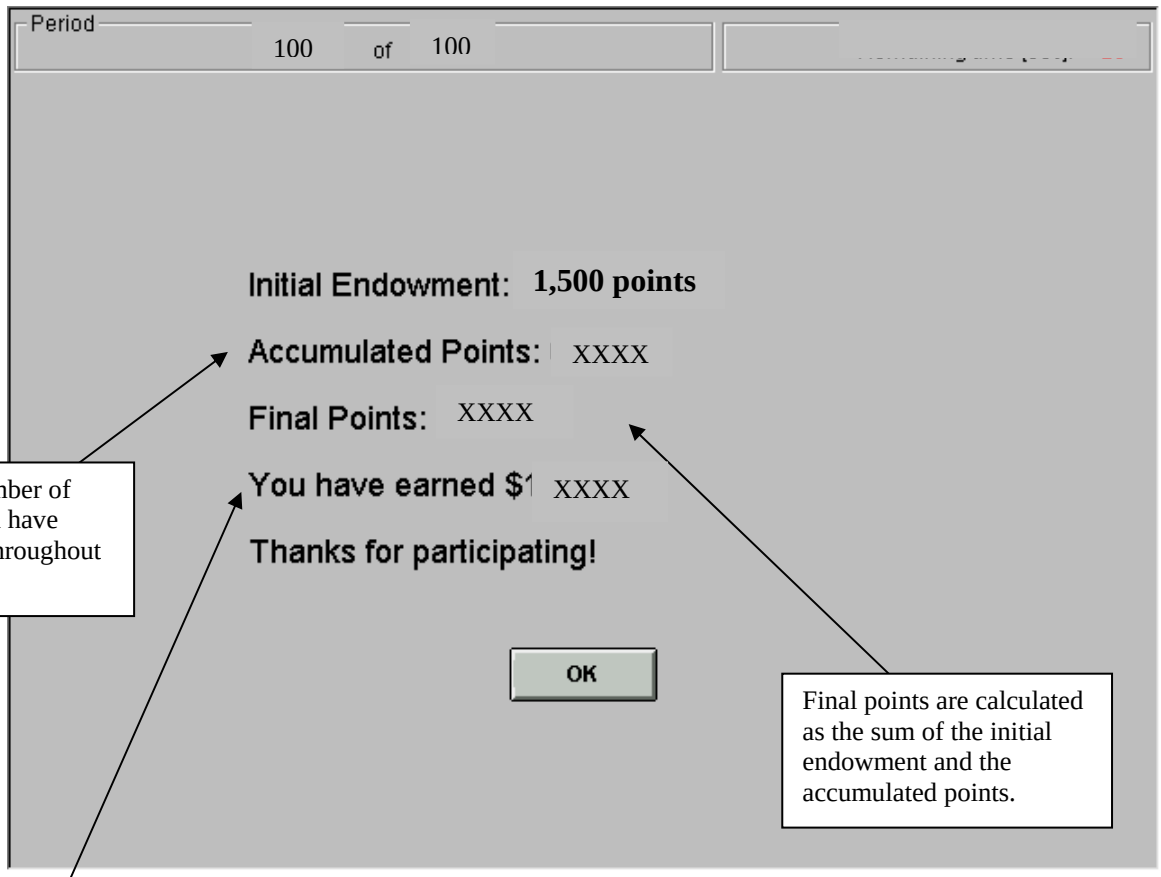

This is the number of points that you have accumulated throughout the 100 trials.

Final points are calculated as the sum of the initial endowment and the accumulated points.

This is the amount that you will receive in cash at the end of the experiment. This value is calculated by applying the following exchange rate: 10 points equal 7 cents.

Brit Grosskopf, Texas A&M University
Yoella Bereby-Meyer, Ben Gurion University
Max Bazerman, Harvard Business School

Corresponding author:

Brit Grosskopf
Department of Economics
Texas A&M University
TAMU 4228
College Station, TX 77843-4228

bgrosskopf@econmail.tamu.edu
Tel: 979-845-7357
Fax: 979-847-8757