The Origin of the Winner's Curse: A Laboratory Study

Gary Charness and Dan Levin[*]

January 24, 2005

Abstract: The Winner's Curse (WC) is one of the most robust and persistent deviations from theoretical predictions that has been established in experimental economics. We designed and implemented in the laboratory a simplified version of the *Acquiring a Company* game. Our transformation reduced the game to an individual-choice problem, but one that still retains the key adverse-selection issue that results in the WC. This allows us to isolate the origins of the WC from proposed explanations that have recently been advanced, such as ignoring others' cognition process, having cursed system of beliefs, high variance in payoffs, and/or a misunderstanding of the game. Our main results find that the WC is alive and well in environments where the above explanations are either mute or wrong, and strongly suggest that the WC is better explained by bounded rationality of the form that people have difficulties either performing Bayesian updating or seeing through a complex problem.

Key Words: winner's curse, auction, experiment, take-over game, acquire-a-company game, self-cognition bias, cursed (equilibrium) beliefs.

Contact: Gary Charness, Department of Economics, University of California, Santa Barbara, 2127 North Hall, Santa Barbara, CA 93106-9210, charness@econ.ucsb.edu, http://www.econ.ucsb.edu/~charness. Dan Levin, Department of Economics, The Ohio State University, 1945 N. High Street, Columbus, OH 43210-1172, dlevin@econ.ohio-state.edu, http://www.econ.ohio-state.edu/levin/

1. Introduction

The *winner's curse* (WC) is a well-known phenomenon in *common-value* (CV) auctions and bidding behavior. Bidders systematically fail to take into account factors in auction environments indicating that the winning bid is very likely to be an overbid, resulting in an expected loss. The two most general and robust results that emerge from past experimental studies of the WC are that the WC is alive and well and that it is persistent, vanishing very slowly if at all. Several explanations, which we review below, were proposed to explain the existence and persistence of the WC. In this paper we examine experimentally these past proposals and contrast them with our own explanation.

We designed an experiment that allows us to isolate and focus on particular explanations: 1. The WC results from a failure to consider the cognition of other bidders. 2. The WC results from a particular and cursed belief structure regarding other players' rationality. 3. The WC is simply a failure to fully take into account the "rules of the game," in that they do not absorb essential features of the structure of the game or environment *per se*, but instead tend to simplify. 4. The WC and its persistence are due to unclear feedback after incorrect actions (decisions) resulting in large variance in payoffs that creates a very noisy and difficult learning environment. We contrast these explanations with our own presented below.

Consider a single unit, *common-value* (CV) auction, where the *ex-post* value of the auctioned item is the same for all bidders. Assume further that each of the $n \geq 2$ bidders' privately observes a signal (an estimate) that is an unbiased estimator for the CV and correlated with it and the other signals. It is reasonable to expect that the bids submitted will be highly correlated with the bidders' estimates and as a result the winner, the bidder with the highest

1

submitted bid, also holds (one of) the highest estimate. Thus, upon winning the winner learns that her *ex-ante* unbiased estimate is in fact (one of) the highest.

Fully rational decision-makers are expected, while making their present decisions, to condition upon the critical future event (e.g., winning the auction) and correctly infer and incorporate the relevant posterior in their current bidding decisions. A systematic failure to do so in CV auctions results in overbidding and systematic negative (or below *normal*) payoffs. This systematic failure represents departure from rationality in individual choice problems and or, in combination with other departures from standard game-theoretic assumptions, in strategic situations.[1]

The first formal claim of the WC was made by Capen, Clap and Campbell (1971), three petroleum engineers, who argue that oil companies had fallen into such trap and thus suffered unexpected low profit rates in the 1960's and 1970's on OCS lease sales "year after year."[2,3] Not surprisingly, many economists greeted this claim with a 'healthy' dose of skepticism. After all, such a claim implies repeated errors and a departure from rationality, or at least from equilibrium predictions.[4] Nevertheless, that paper sparked a small literature of investigators to estimate rate of returns from accounting data.[5] The empirical literature does not reach a clear and unambiguous conclusion about the existence of the WC, perhaps because there are many obstacles that complicate enormously a study based on "real life" data. For example, studying

---

[1] Unfortunately, in discussion of the WC economists often use the term to refer to the difference between the expected value of the item conditional on the event of winning and the "naïve," unconditional on winning, expectation of the CV. Such difference may serve as a measure to the extent of the adverse selection in equilibrium.

[2] Of course Groucho Marx's statement: "I wouldn't want to belong to any club that would accept me as a member" can be viewed as an earlier recognition of the WC.

[3] Capen *et al.* conclude that "He who bids on a parcel what he thinks it is worth, will, in the long run, be taken for a cleaning."

[4] For such a claim made by others see Lorenz and Dougherty (1983) and references therein.

[5] For example see Mead, Moseidjord and Sorensen (1983, 1984), Hendricks, Porter and Boudreau (1987), and Leitzinger and Stiglitz (1984).

the existence of the WC merely by estimating rate of returns does not control for several key economic variable that affect behavior in such environment (e.g., bidders' private estimates).

Laboratory experiments are an alternative approach for studying the WC. The first well-known experiment in this realm is the classroom "Jar Experiment," conducted by Bazerman and Samuelson (1983), where a jar of coins worth $8.00 were auctioned off to a class, with prizes for the best guess. While the average estimate of the jar's value was $5.13, the mean winning bid was nevertheless an overbid, $10.01. One important characteristic that distinguishes this experiment from most of the others that followed is that neither the distribution functions of the common value nor that of private estimates was induced. In other words, the informational structure was not controlled and at least formally subjects/bidders observed the same public information, a jar full of coins of different nomination. With such an omission it is impossible to solve for and calculate departures from the risk-neutral Nash equilibrium benchmark. It is also impossible to find out much about individual departures since we cannot observe the (private-estimate) information upon which they based their bids.[6]

Partially to overcome this omission, Kagel and Levin (1986), and Kagel and Levin in collaboration with others, study the existence of the WC and closely related issues resulting from such existence in a sequence of later papers. The main finding is that the WC is "alive and well," and is a robust phenomenon in many CV auction forms, such as the first-price auction, the second-price auction (Kagel, Levin, and Harstad 1995), and the English auction (Levin, Kagel, and Richard 1996).[7] Clearly, such persistent losses (or below-normal profits) are not part of any equilibrium behavior with fully rational bidders. In addition, Dyer, Kagel, and Levin (1987)

---

[6] For example, it is not clear if overbidding by winners resulted from a failure to incorporate order-statistic calculations or from an extreme visual perception bias on the part of the over bidders (who perhaps did correct for the usual symmetric adverse-selection problem).
[7] See also Kagel and Levin (1991) and Kagel and Levin (1999).

report an experimental auction with a group of "experts" (managers of construction firms); these presumably more sophisticated agents did no better or worse than students.

It thus becomes important to study if such departures from equilibrium and/or rationality are eliminated (mitigated) over time or with subjects with experience from previous sessions. The finding here both in auction environments (Kagel and Levin 2002) and in the "Acquiring a Company" task (discussed in detail below) that were studied in depth are that: 1) Learning to cope with the WC is very slow and though "worst offense" behavior is at least mitigated, behavior does not seem to converge to the relevant equilibrium, and 2) To the extent that there is learning, results suggests that it is situation-specific and does not generalize since losses re-emerge with a treatment that worsens the adverse-selection problem (e.g., more bidders). Documenting such disequilibrium results has significant implications for public policy. For example, releasing additional public information is predicted to raise seller's revenue. However, in auctions with the WC such additional information seems to be used to calibrate and discipline over-optimism and hence results in lower bidding and lower revenues.

In fact, the WC is a phenomenon found in a variety of environments other than auctions. Cassing and Douglas (1980) and Blecherman and Camerer (1998) find it in baseball's free agency markets. Dessauer (1981) suggests it exists in bidding for book publishing rights. Roll (1986) suggests it exists in corporate takeovers. Rock (1986) presents a model in which the observed underpricing of initial public offerings is a direct result of the winner's curse problem facing uninformed investors, and Levis (1990) provides empirical evidence (from British IPOs) that suggests that the winner's curse and interest rate costs explain this underpricing.

Experimental work related to the WC has become a small literature, surveyed in Kagel and Levin (2002). It seems that presently only a few economists doubt the existence of WC-type

of behavior in the lab and more and more are willing to accept that such behavior may indeed exists outside the lab in real markets.[8]

Given the main results above it is important to understand the origin of such departures and identified decision-making environments that either exacerbate the problem or ease it. We describe and discuss some stories/explanations for the WC in Section 2, and present our experimental design objectives and predictions in Section 3. Section 4 contains our experimental results, which are discussed in detail in Section 5; Section 6 concludes.

2. The Origin of the Winner's Curse

What are the underlying determinants of the WC? For one reason or another, competitive decision-makers fail to properly incorporate all of the relevant information when choosing an action. An important question is whether we can identify conditions or methods that tend to ameliorate this problem, whether it is caused by cognitive limitation or some form of psychological bias.

Optimal behavior in a multi-person auction involves complicated calculations of one's best response, involving both beliefs about rivals' rationality and strategic uncertainty. Thus, it might be too much to expect that Nash equilibrium prediction will organize experimental data well. Furthermore, given the complexity of the environment it is difficult to clearly ascertain the cause(s) of departures from the Nash norm. It is therefore useful to test behavior in a simpler environment where we can observe an individual's decision process in isolation, rather than as

---

[8] Obviously, there are some exceptions. A notable one is Cox, Dinkin, and Swarthout (2001), who argue that free entry and exit will cure the problem. In their work they show that only a subset of potential bidders chose to participate in such a CV auction (participation was costly), while the worse offenders choose to stay out. This resulted in CV auctions with a lower number of bidders where the WC was largely mitigated. However, one may object to the line of reasoning by which this process diminishes the WC on the grounds that an epidemic that wipes out the infected population is not exactly a cure.

part of a more complex interaction involving other parties. Such a simpler environment may help isolate origins of the WC that cannot be explained by failure of beliefs in common knowledge of rationality and/or best-response considerations.

One such environment is found in the "Acquiring a Company" game (also known as the "take-over game"), which is based on the famous "lemons" problem in Akerlof (1970) and was first described in Samuelson and Bazerman (1985). In spite the simplicity of this game, which abstracts from many of the complications embodied in an auction context, it is critical to emphasize that overcoming the WC – the adverse-selection problem – is still the key difficulty. The game has two players, a bidder (firm) and a responder (the target firm). The bidder makes a "take-it-or-leave-it" bid and the responder either accepts the bid or rejects it; in either case the game is over. The bidder is faced with the task of making an offer for a firm with a current value known to the firm, but unknown to the bidder. The bidder does know that the current value is an integer between 0 and 100, inclusive, with each value equally likely; the bidder also knows that the firm's value will increase by 50% if acquired.[9] The responder transfers the firm to the bidder if and only if the bid is at least as high as the value to the current owner.

Upon suitable reflection, the bidder should realize that whenever the responder accepts the bid, B, it implies that the owner's value does not exceed the bid. This conclusion eliminates the possibility of high values (the adverse-selection issue) and necessitates forming a new and relevant expectation regarding the firm's value, incorporating the posterior distribution. With a *uniform* original distribution the posterior is also a uniform distribution on [0, B]; thus, its expected value is $\frac{B}{2}$. This leads to the conclusion that any non-negative bid (smaller than 100)

---

[9] The domain of values in the original game is the continuous interval [0,100]. However, we simplify by casting   the description closer to our actual design.

yields an expected value of $\frac{3}{2} \times \frac{B}{2} = \frac{3B}{4}$, for an expected loss of $\frac{B}{4}$, and at this point it is clear that the optimal bid is zero.

Despite the simplicity of this argument, in practice the bid chosen is typically somewhere between the unconditional expected value to the seller of 50 and the unconditional expected value to the acquirer of 75.[10] Such bidding behavior that generates systematic losses, the WC, has been found to be ubiquitous in laboratory tests despite a variety of interventions designed to overcome it. Ball, Bazerman, and Carroll (1991) find that few participants overcome the WC even when buyer and seller roles are reversed over the course of the session, done in an attempt to highlight the informational asymmetry. Selten, Abbink, and Cox (2001) find similar bidding behavior when the optimal response is to bid a positive amount and there are 100 trials in a session.

Bereby-Meyer and Grosskopf (2002) introduce a variation in which the participants were given an option to not bid, as there is a concern that participants otherwise bid positive amounts simply because they assume that they are supposed to do so; the median bid is 35. In another condition, they reduce the payoff variability by requiring participants to submit a bid that applies to each of 10 firms.[11] Feedback was given about the average profit made; in one condition the value of each of the firms and the profit from each individual outcome was also reported. This led to some success in reducing the WC, as 18 of 54 participants (in the two conditions) bid close to zero in at least the last three periods of the session. Bereby-Meyer, Grosskopf, and Bazerman (2003) try to "enhance the decision-makers' cognitive understanding of the task" by varying

---

[10] Ball, Bazerman, and Carroll (1991) and Tor and Bazerman (2003) also use a protocol where they ask subjects to write down their reasoning. A typical explanation resembles: "Well, on the average the value is 50 which means it is 75 for me so if I offer (say) 65 we both make handsome profits."

parameters and even giving participants feedback about foregone payoffs, but none of these approaches generated a better understanding of the task.

In light of these results, it seems fair to say that the WC is robust to a variety of interventions in the take-over game, in contrast to early criticisms that it was merely a transient and local phenomenon.

A number of explanations for the existence and stubborn persistence of WC phenomena have been proposed in recent years. Many of these involve the idea that individuals tend to focus to their own cognition while ignoring or minimizing a similar process by another person engaged in the decision task whose action ought to impact one's own decision. For example, Carroll, Bazerman, and Maury (1988) and Tor and Bazerman (2003) find that buyers seemed to treat the decision as if there was no asymmetry of information, ignoring the sellers' choice process. Furthermore, Bereby-Meyer, Grosskopf, and Bazerman (2003) find that bidding behavior in a symmetric-information treatment is similar to that in the standard asymmetric-information treatment.

The Eyster and Rabin (2002) concept of the *Cursed Equilibrium* provides an elegant formalization of the psychological principle that one tends to under-appreciate the cognition processes and/or the informational content of others. They define and apply this notion to a variety of economic environments ranging from common-value auctions, Akerlof's lemon market, the take-over game, and voting in juries. Their cursed equilibrium is an attempt to address the overwhelming (and growing) evidence, coming in particular from experimental data, that game-theoretic predictions based on common-knowledge and fully rational decision-makers often miss the mark completely. Their model players who do not fully take into account how

---

[11] A similar aggregation device, requiring investors to choose investments for three periods, but reporting only the

other people's actions are contingent on these others' private information.[12] As we explain below, such imperfect accounting may help explain systematic overbidding; this is the WC in both common-value auctions and the take-over game.

In the cursed equilibrium, "each player correctly predicts the distribution of other players' actions, but underestimates the degree to which these actions are correlated with these other players' information." Applied (for simplicity) to the continuous *uniform* [0,100] version of the take-over game a (fully)[13] cursed bidder correctly recognizes that by offering any payment, B, such that B > 0, a fraction (probability) of B/100 (only those types with firm values that are smaller than B) will transfer the company. However, the cursed bidder then ignores the specific types that will (refuse to) transfer the firm and applies the probability B/100, non-discriminatively to any type. Solving the bidder maximization problem under such cursed beliefs yield the following problem:

$$Max_{B \geq 0} \int_{0}^{100} \left[\tfrac{3v}{2} - B \right] (B/100) f(v) dv$$

The solution to this problem is B*= 37.5. And although 37.5 is still significantly lower than observed data in the earlier Bazerman *et al* experiments (in fact, the 37.5 prediction comes quite close to some of our data), the cursed equilibrium close the gap between theory and evidence. However, if we strip the 'other' person from the game (converting it to an individual-

---

outcome aggregated over the three periods was found by Gneezy and Potters (1997) to be effective in reducing myopic loss aversion.

[12] Eyster and Rabin formalize and generalize more   *ad hoc* attempts by Kagel and Levin (1986) and Holt and Sherman (1994) to explain deviation from game-theoretic predictions based on full rationality.

[13] Eyster and Rabin model connects the spectrum from a fully rational player with to a perfectly cursed player by introducing an additional parameter λ. For example, $\lambda = 0$ and $\lambda = 1$ respectively mean that the player is fully rational or perfectly cursed in his or her beliefs. Intermediate values such as $\lambda = 1/2$ mean that the player behaves as if 1/2 are fully cursed and 1/2 are fully rational. Here even $\lambda = 1$ does not move predictions enough so we omit discussing it in the text for brevity.

choice problem) and still observe almost the same extent of the WC, the cursed-equilibrium story would not apply to the origin of the winner's curse, at least with respect to the take-over game.

While we agree that it is often critical to recognize the cognitive processes of others, we demonstrate in this work that the origin of the WC is not in a bidder neglecting the cognitive processes of other parties. Rather, the WC stems from a decision maker inability to recognize that a 'future' event *per se* is informative and relevant for their current decisions and/or an inability to correctly process such information and construct the correct posterior.[14] Thus, the complexity of the environment may be an important factor, as previous studies (e.g., Charness and Levin forthcoming) have shown that while people may be poor Bayesians in a complex environment, they do quite well in a very simple one.

Bereby-Meyer and Grosskopf (2002) interpret their results as indicating that the larger is the variance in the feedback, the slower is the learning. In a single take-over game, one earns a positive payoff 1/3 of the time with any positive bid. On the other hand, when the same bid is used for 10 separate draws and the results are aggregated, the payoff variance is indeed greatly reduced. The argument goes that variability in the environment impedes adaptive learning, consistent with the notion that partial reinforcement slows learning.

Although these arguments are quite reasonable, we conjecture that the WC instead stems from a failure to condition one's bid on the information provided by the final realization, compounded by poor updating when this idea is even considered. Accordingly, we design an experimental test that totally strips away any element of another party's decision process; in short, we "eliminate the middle man".

---

[14] Tor and Bazerman (2002) argue that "the rules of the game and the decisions of other parties are pieces of information that are typically overlooked" by decision-makers, leading to suboptimal behavior. While we think that this may indeed be true in some cases, we don't think that our bidders fail to understand the simple rules

We are not the first to use a computer to interact with a bidder: Bereby-Meyer and Grosskopf (2002) and Grosskopf, Bereby-Meyer, and Bazerman (2003) report treatments in which a computer acted as the seller of a company. However, their designs explicitly framed the decision as purchasing a company from a firm and encouraged bidders to think of this as a transaction between two parties. In contrast (as becomes apparent in the next section), we make no mention of any other party, and it is completely transparent to the bidder that there is no other entity involved, just one's bid and one's choice of card on the screen.

Our simple design enables us to examine and contrast our conjecture to the alternatives mentioned. We also conduct some additional treatments with a restricted set of possible values, with the feature that the payoff variance is much higher but the complexity of the environment is much lower.

3. Experimental Design and Predictions

We conducted 12 experimental sessions (three sessions for each of four treatments) at the University of California at Santa Barbara. Participants were recruited from the general campus population by an e-mail message; 206 people participated in our sessions. We distributed written instructions to the students and read these aloud, taking questions as they arose; these instructions are presented in Appendix A. The experiment itself was performed using a web-based computer interface.

As mentioned above, one of our key design objectives is to test whether it is a tendency to neglect the thinking processes of other parties and/or a cursed belief system about rivals that drives the winner's curse, or whether overbidding stems from an inability to reason through the

---

determining the outcome of their bids. Instead, the "rule" not understood here is the need to update the support of

problem even when there are no other decision-makers involved, due to: 1) a limited ability to condition on the additional information embodied in a 'successful bid' and incorporate it in current decisions, 2) a limited ability to use Bayesian updating rule, and/or 3) a limited ability to overcome the complexity of the task at hand without respect to the involvement of other actors. Accordingly, our design neither has a responding seller nor even mentions such a role. In other words we have transformed the take-over game to an individual-choice task. After the instructions were read and discussed, 100 'cards' were displayed on each person's computer screen, in a 10x10 array. Each card had an integer value between 0 and 99, inclusive, with one card having each of the possible 100 values. In the 'standard' version, the cards were initially displayed in sequence (0 to 9 on the first row, 10 to 19 on the second row, etc.), so it was easy to verify that there was exactly one of each of these integers; in the other versions, monotonicity was also preserved in the initial presentation.

After viewing these cards, a participant clicked to flip the cards over and start the game, knowing that the cards would be re-shuffled after every action (as he or she could verify during the session). Each period the decider chose a bid (a non-negative integer) and then selected one of the cards; the card's value was then revealed. If the bid was at least as large as the value displayed, the participant paid the bid and received 1.5 times the value displayed on the card; if the bid was (strictly) less than the value displayed, there was no bid payment or value received. This outcome was displayed on the screen.

Clearly there are no external cognitive processes at issue here, so that any overbidding cannot be attributed to ignoring others' cognition or having cursed beliefs. Instead, systematic overbidding must be the result of other forms of bounded rationality, as hypothesized below.

---

the relevant distribution of value.

We also utilized two other distributions of value, designed to shed light on whether people are simply bad Bayesians, and have difficulty updating their priors when given new relevant information, or whether it is the problem of high payoff variance that allows the winner's curse to persist over time. We reduce the complexity of the Bayesian updating problem while simultaneously increasing the variance of the outcome and feedback.

In one version a card can have only two possible values, either 0 or 99. A participant sees 50 cards with a value of 0 and the other 50 cards with a value of 99. Thus, the only 'reasonable' bids are either 0 or 99,[15] as any positive bid less than 99 is strictly dominated by a lower bid: A higher bid (less than 99) does not increase the probability of a successful bid, but does raises the cost to obtain it. Note that this is indeed a vast simplification as it involves only *ordinal* logic. The only element of *cardinality* in the process is to compare bidding 0 to bidding 99; since bidding 0 yields a certain payoff of zero, while bidding 99 produces a rather unattractive lottery where the bidder wins 49.5 or loses 99 with equal probabilities of 1/2. Nevertheless, even in this case the two-value treatment (hereafter the "2VT") with its much simpler logic allows one to avoid using formal Bayesian updating to form posterior expectations. On the other hand, note also that the variance of the value is much greater in the 2VT (2450.25 vs. 833.25).[16]

In the 100-value treatment (the original game, hereafter the "100VT"), the analysis is much more complex. Why is a bid of 60 worse than a bid of 59? If the firm's value is 61 or more it does not matter. If the firm's value is precisely 60, the bidder is better off by 30 bidding 60 and winning than unsuccessfully bidding 59. In all of the other 60 cases the bidder is worse

---

[15] Or 100 if one is uncertain if the bid needs to exceed the value or just equal it.

[16] In the original game a bidder makes a profit conditional on winning whenever the firm's actual value is between the bid and two thirds of the bid (a probability of 1/3) and make a loss whenever the firm's actual value is between 0 and two-thirds of the bid (a probability of 2/3).

off with a bid 60 than a bid of 59, but only by 1. Thus, recognizing that bidding 59 dominates (in expectations) bidding 60 requires some non-trivial cardinal reasoning.

After observing our outcomes in these treatments, we also conducted an intermediate four-value treatment ("4VT"), in which there were 25 cards each with values 0, 33, 66, and 99. Obviously, this is less complex than the 100VT, but is nevertheless qualitatively different than the 2VT and could provide further insights into the nature of the WC in our setting. We chose the 4VT in part to examine the degree to which the increased proportion of zero bids observed in the 2VT was an artifact of there being only two 'reasonable' bids.

We note that bidders were free *in all treatments* to choose any non-negative integer as a bid. While we shall see that bids in the restricted-values treatments tend to cluster at the feasible values, participants made such choices without guidance.

Each of our experimental sessions consisted of 60 periods. In three of the sessions, participants were first presented with cards having 100 distinct values (henceforth we refer to this treatment as 100VT). They were told that the distribution of the cards' values would change after 30 periods, and the new values would be displayed at that time. In the subsequent 30 periods (henceforth we refer to this treatment as 2VT), cards had only two possible values {0, 99} as discussed above. In three other sessions, participants were first presented with the 2VT, and the 100VT began after 30 periods. In a third set of three sessions, there were four segments of 15 periods each, in the order 2VT-100VT-2VT-100VT. In our final set of sessions, we had two 30-period segments, in the order 4VT-100VT. In all cases, a history of the results of all previous periods in a segment (showing period, bid, value, profit) was displayed on the screen.

Participants received $5 for showing up on time. Since we expected losses, they also received a $15 endowment. At the end of the session, we randomly chose one period from each

14

10-period block for payment, so that only six periods actually counted towards monetary payoffs. This device helped us to avoid (or at least minimize) possible intra-session income effects that could readily affect behavior. After netting out the experimental gains and/or losses from the six chosen periods, we converted these to actual dollars at the rate of $0.10 for each unit, and added the result to the original endowment. While one could lose the entire $15 endowment, no participant could leave with less than the $5 show-up fee.

We have two conjectures concerning our results. First, we strongly suspect that the winner's curse will persist within our new structure of the take-over game even without a seller, so that the argument that this behavior results from an inability to consider the cognitive processes of others or cursed beliefs about others would appear to miss the point. Second, we explain below the sense in by which the 2VT significantly simplified the task and reduced the complexity. Thus, we expect that it will help ameliorate the winner's curse, even though the variance is higher.

4. Results

We find that the winner's curse is alive and well when there is no seller. We also observe very different patterns of bidding behavior depending on the number of distinct values for the cards. Table 1 summarizes our overall results:

Table 1: Aggregate Bidding Statistics, by Sequence in Session

Session Sequence	# obs., each VT	Avg. bid, 100VT	Avg. bid, 2VT or 4VT	Zero-bid rate, 100VT	Zero-bid rate, 2VT or 4VT	Exp. loss, 100VT*	Exp. loss, 2VT or 4VT*
100VT, 2VT	1620	38.86	61.41	.0753	.3346	9.72	16.09
2VT, 100VT	1440	37.01	75.51	.2090	.3042	8.98	15.48
2VT, 100VT twice	1410	33.61	41.14	.1823	.4929	8.37	10.51
4VT, 100VT	1710	32.64	37.69	.3099	.3058	8.08	7.23

*We consider only bids not greater than 100 in making this calculation.

At a first glance, we see that the majority of bids were positive in all cases. However, around 30-50% of all bids are zero when we limit the number of possible values to two or four, so that we do see significant amelioration of the winner's curse in the less complex environment. Interestingly, there appears to be learning from being exposed to the limited-value case, as zero bids in the 100VT are much more frequent when it followed a 2VT or 4VT (21%, 18% and 30.5%, as compared to 7.5%). Note that the average bid is substantially higher in the 2VT than in the 100VT (although less so when there are four segments), so that participants actually lose more money with the simpler design. On the other hand, the average bid in the 4VT is nearly as low as the average bid in the 100VT.

Expected losses, given the actual bid distribution less any bids over 100, were always higher in the 2VT than in the 100VT. Thus, using average losses as a benchmark, one might be tempted to conclude that the winner's curse gets worse instead of better when there are only two possible values. However, it is instructive to consider the expected losses of bidders in the two cases. In the 100VT, if a bidder picks a random integer from [0,99], the expected loss is

$\sum_{0}^{99} \frac{j+1}{100} \times \frac{j}{4} \times \frac{1}{99} = 8.33$. In the 2VT, if a bidder picks a random integer from [0,99], the expected

loss is about 24.20, nearly three times as great.[17] In this sense, bidders are doing far better than

random in the 2VT (actual losses of 14.03), but not in the 100VT (actual losses of 8.8 overall).

Note that even if a bidder in the 2VT is somewhat more sophisticated, choosing zero half of the

time and 99 half of the time, the expected loss is 12.38, nearly half as much again as the expected

losses with random bidding in the 100VT.

Finally, we point out that bidding is clearly improved in the 4VT sessions, as both the

zero-bid rate is higher in the 4VT than in the 100VT and the expected loss is lower in the 4VT

than in the 100VT.

Figures 1-3 show the bid frequencies that support these results:

Figure 1: Bid Frequency in 100-value Treatment

---
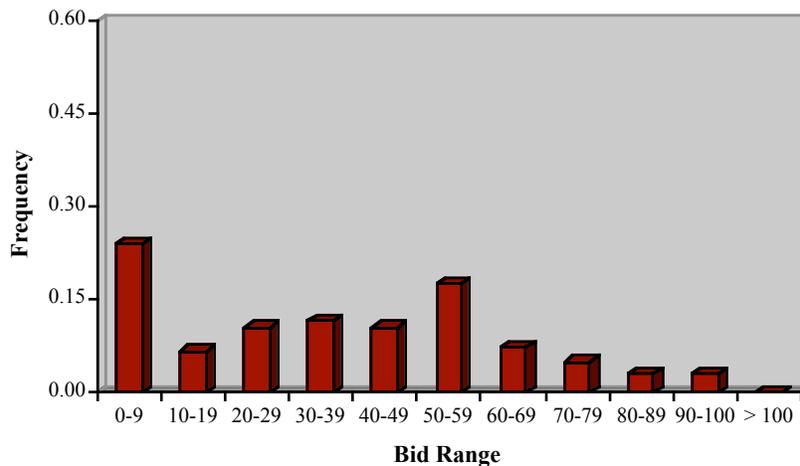
[17] We note that there was nothing in the instructions that would preclude bidding values other than 0 and 99.

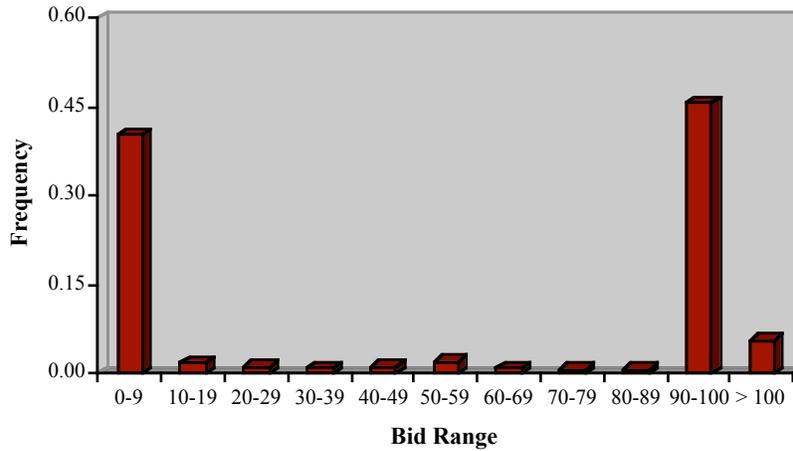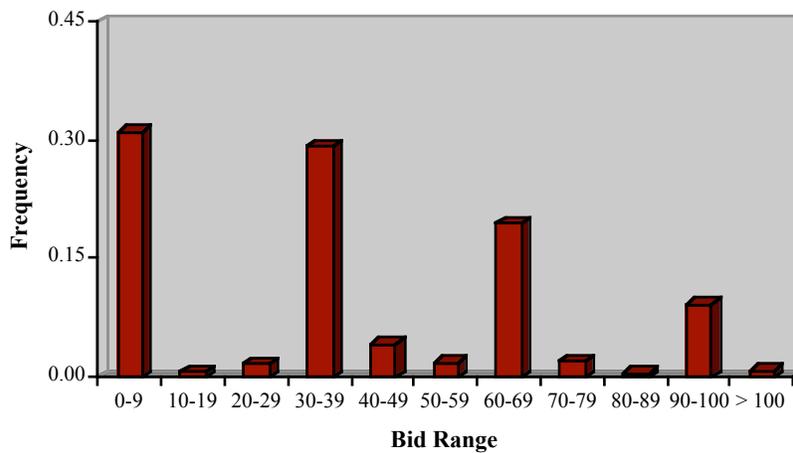**Figure 2: Bid Frequency in Two-value Treatment**


Figure 3: Bid Frequency in Four-value Treatment

The distribution of bids looks fairly normal in the 100VT, with the exception that 24% of all bids were less than 10 in this case; otherwise the most common bid range was 50-59.[18] In comparison, the distribution of bids in the 2VT is heavily bi-modal, with either very low bids or bids near 100 predominating; only about 10% of all two-value bids were between 10 and 90. The pattern has four modes in the four-value treatment, one at each of the values. Thus, it seems

---

[18] Quite in line with what emerges from the Ball *et al.* (1991) protocols.

that participants were mainly able to grasp the notion that there was no sense in bidding anything other than one of the values.[19]

A natural consideration is whether there are trends in behavior over the course of a 30-period segment. In Figure 4, we consider the zero-bid rate over time for 5-period ranges of the 100VT in each type of session:

**Figure 4: Zero-bid Rate over Time in 100VT segments**
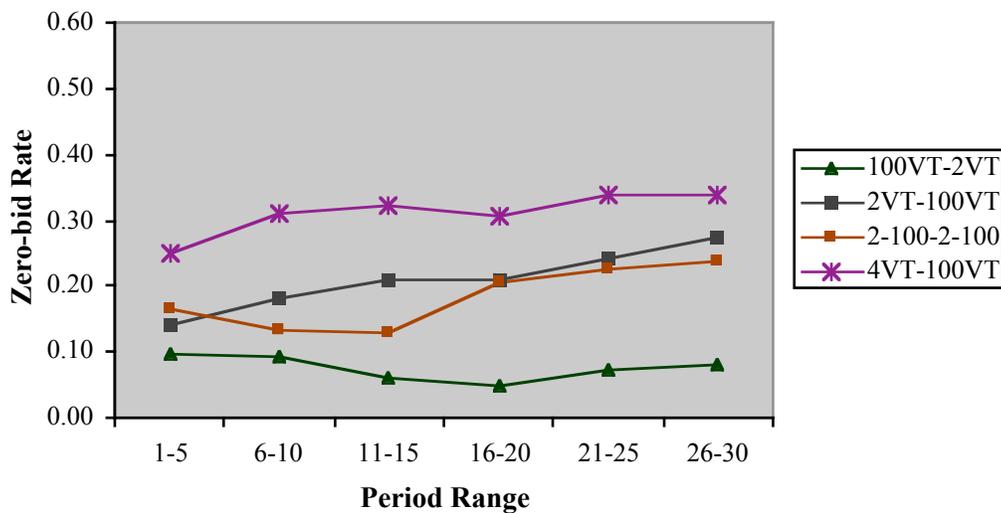
As is almost invariably the case in studies where there are 100, or a continuum of, possible values, the winner's curse (here all positive bids) does *not* diminish over time in our 100VT *when there has been no exposure to the 2VT or the 4VT*; if anything, the zero-bid rate declines over time. However, the pattern in the 100VT is different when it follows either the 2VT or the 4VT during a session, as the zero-bid rate increases over time in each of these cases. The zero-bid rate is also much higher in the 100 VT when it follows a limited-value treatment

---

[19] People in the two-value case bid zero (37.5%) or ninety-nine (37.5%) three-quarters of the time; if we also include bids equal to 100, this comprises 82.5% of all bids. Similarly, in the four-value treatment, zero was bid 30.6% of the time, 33 was bid 23.1% of the time, 66 was bid 15.1% of the time, and 99 was bid 7.0% of the time; these four bids comprised 75.7% of all bids (85.0% if we include bids equal to 34, 67, or 100).

than when it doesn't: in the final five-period segment, the corresponding rates for the four cases are 8%, 28%, 24%, and 34%. Clearly the exposure to limited values improves the zero-bid rate; it is interesting that much of this effect seems to take place over the course of the 100VT.

In Figure 5, we consider the zero-bid rate over time for 5-period ranges of the 2VT or 4VT in each type of session:

Figure 5: Zero-bid Rate over Time in 2VT & 4VT segments
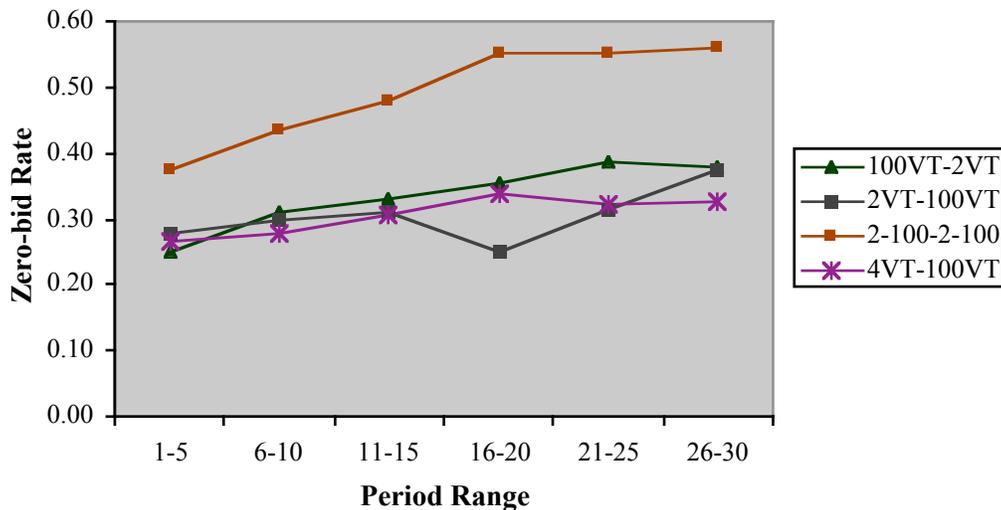
Here we see a positive trend in the zero-bid rate in all cases, with the most dramatic gains over time in the four-segment treatment and the least change in the 4VT. Perhaps the large variance in outcomes when bidding 99 in the 2VT serves to enhance learning, while having the less extreme 'reasonable' positive bids of 33 and 66 in the 4VT slows this process down.

The data we have shown are aggregations; however, this presentation masks the fact that there is considerable heterogeneity in behavior, particularly in the 2VT. To illustrate this heterogeneity, Figures 6 and 7 show the number of zero bids made by each individual in the 100VT or the 2VT and 4VT:

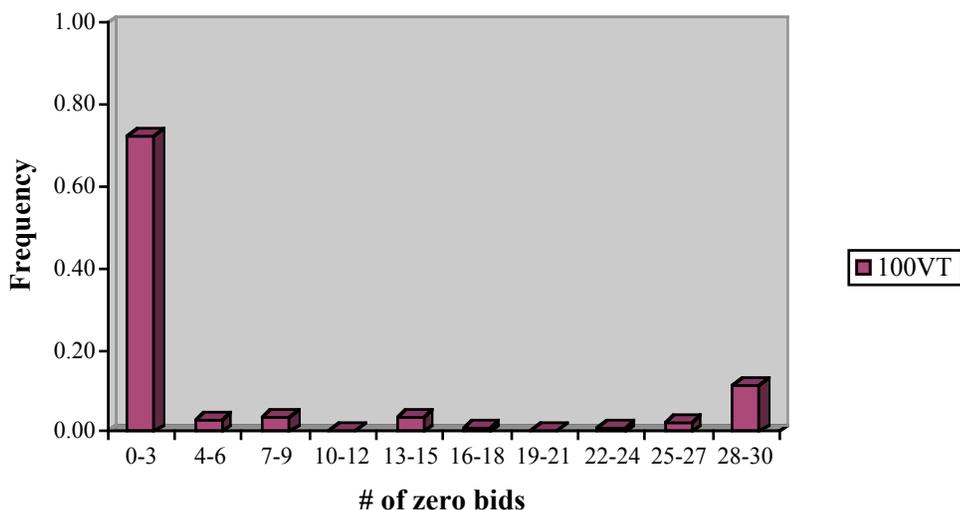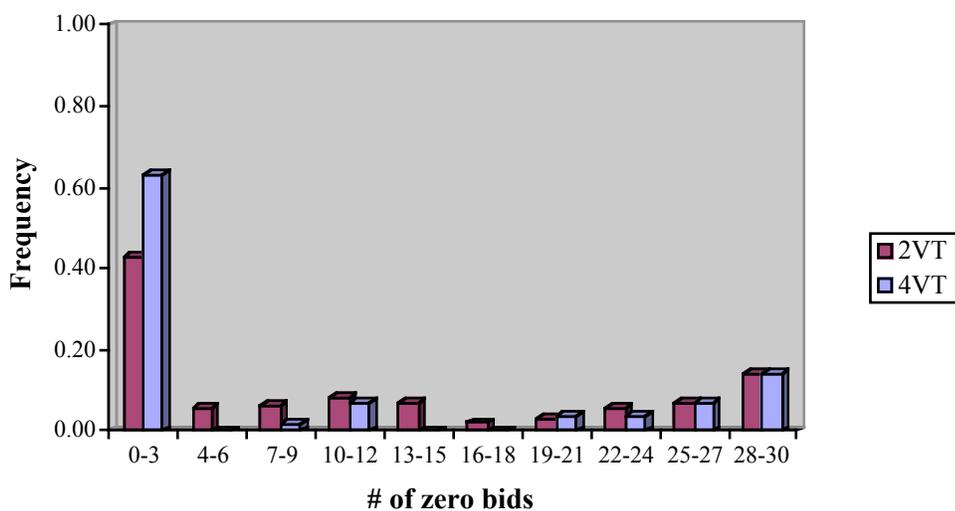**Figure 6: Zero Bids in 100VT, by Individual**

**Figure 7: Zero Bids in 2VT and 4VT, by Individual**

We see that the great majority of the participants rarely bid zero in the 100VT; 79% of the participants bid zero less than 1/3 of the time, while 15% of the participants bid zero more than 2/3 of the time. In the 2VT and 4VT, 57% of the participants bid zero less than 1/3 of the

time, while 27% bid zero more than 2/3 of the time. Thus, the analysis of individual bidding behavior confirms the patterns found in the aggregated data.

In Charness and Levin (forthcoming), we found a difference in behavior across gender in an individual decision-making task involving Bayesian updating. In the present study, we again find a difference in bidding behavior across gender. Overall, the proportion of zero-bids for males is .366, compared to .212 for females; similarly, the average male bid was 36.8, while the average female bid was 49.9. This pattern holds for each of the 2VT, 4VT, and 100VT cases.[20] This result is illustrated in Figures 8-10:

Figure 8: Bid Frequency in 100-value Treatment, by Gender

---
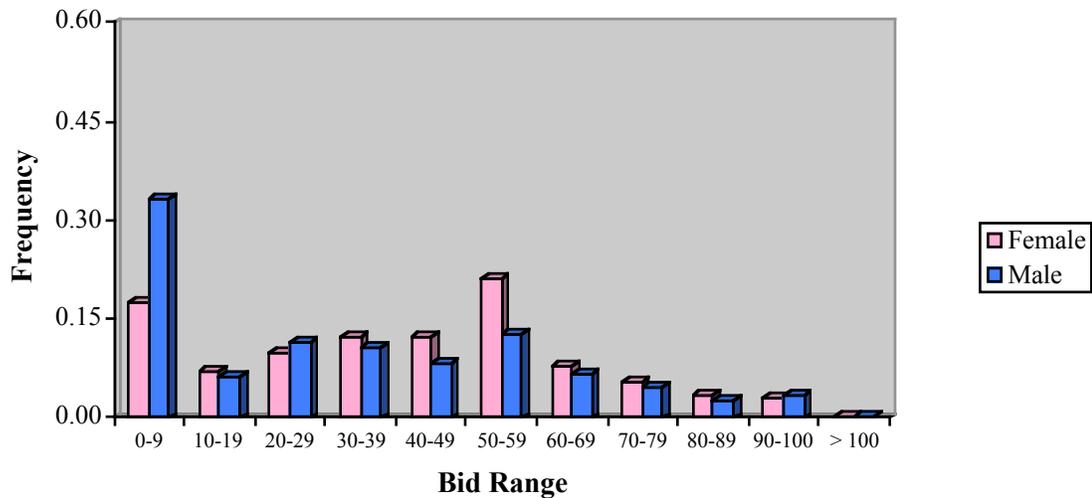
[20] In the 2VT, the zero-bid rate for males was .486 and the female rate was .300; the average bid was 46.2 and 68.6, respectively. In the 4VT, the zero-bid rate for males was .368 and the female rate was .257; the average bid was 35.4 and 39.4, respectively. In the 100VT, the zero-bid rate for males was .282 and the female rate was .135; the average bid was 30.6 and 39.0, respectively.

Figure 9: Bid Frequency in 2VT, by Gender

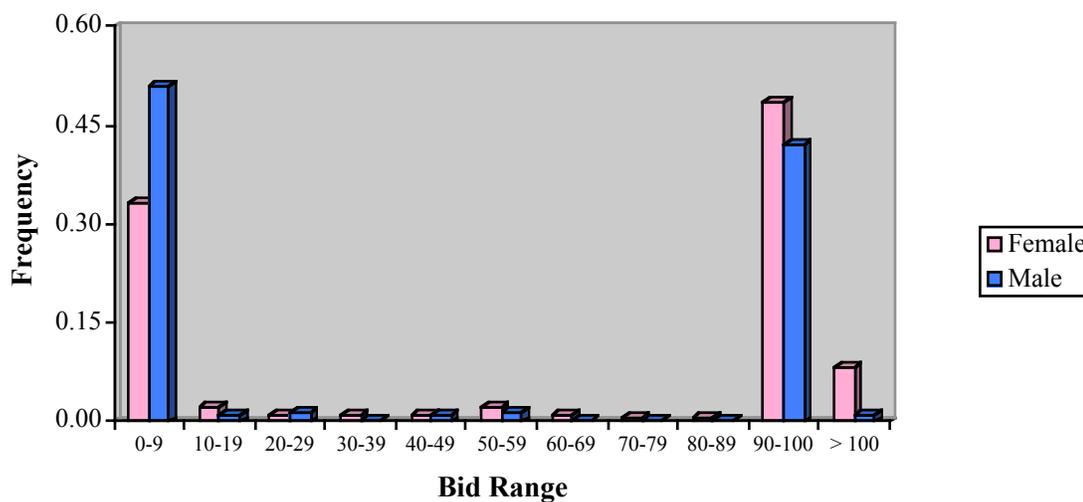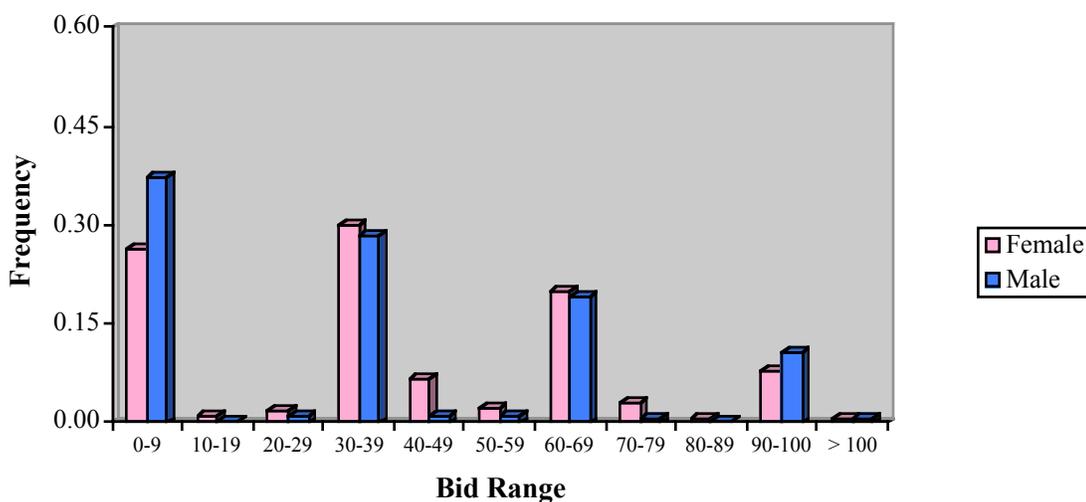**Figure 10: Bid Frequency in 4VT, by Gender**

We next perform a series of regressions to supplement the summary statistics and nonparametric statistical analysis in the foregoing subsection. We test three specifications with zero bid as the binary dependent variable and three specifications with the bid made as the (censored) dependent variable. The zero-bid specifications are probit regressions with robust

standard errors and the bid specifications are random-effects OLS regressions.[21] In one

specification, we consider the overall time trend over the course of a session. In the second

specification, we have separate period dummies for each treatment, with the baseline being the

100VT when it is first in a session. In the third specification, we also include how well

participants did on some questions involving Bayesian updating. These questions (shown in

Appendix B) are an attempt to substitute for the usual demographic questions about a student's

major, SAT score, etc. As our task bears similarity to these updating questions, we hope to

clarify whether any observed gender effect is simply an artifact of a difference in updating skills

across gender, or whether such an effect persists in any case.

---

[21] Results for the more technically-correct tobit regressions are very similar to those for the OLS regressions; we report the latter here for ease of interpretation.

Table 2: Regressions for Determinants of Bidding Behavior

Dependent variable

Independent variables	(1) Zero-bid	(2) Zero-bid	(3) Zero-bid	(4) Bid made	(5) Bid made	(6) Bid made
Constant	-0.905***	-0.874***	-1.172***	47.49***	42.48***	46.87***
	(.184)	(.182)	(.221)	(3.37)	(4.18)	(3.66)
2VT	0.209	0.114	0.275	18.37***	25.58***	20.67***
	(.176)	(.183)	(.212)	(2.36)	(3.93)	(2.55)
100VT	-0.352***	-0.360***	-0.241*	-4.71**	0.227	0.416
	(.127)	(.136)	(.136)	(2.02)	(3.62)	(2.16)
Male	0.480***	0.518***	0.439**	-12.85***	-13.01***	-10.31**
	(.148)	(.147)	(.186)	(4.17)	(4.17)	(4.20)
Period in session	0.011***	0.008**	0.008**	-0.127**	0.147	0.156
	(.002)	(.004)	(.004)	(0.063)	(0.169)	(0.097)
Period*2 VT	-	0.006	0.006	-	-0.391**	-0.519***
		(.005)	(.006)		(0.199)	(0.128)
Period*100VT first	-	-0.036***	-0.041**	-	-0.091	-0.526***
		(.012)	(.020)		(0.213)	(0.151)
Period*100VT not first	-	0.010**	0.007	-	-0.326*	-0.338***
		(.004)	(.005)		(0.194)	(0.117)
# Questions correct	-	-	0.324***	-	-	-7.20***
			(.096)			(2.06)
N	12360	12360	7800	12360	12360	7800
R^2	.059	.078	.111	.036	.038	.081

Standard errors are in parentheses; ***, **, and * indicate significance at $p = 0.01$, $p = 0.05$, and $p = 0.10$ (all two-tailed tests), respectively; 100VT = 1 if the cards had 100 values and is 0 otherwise; 100VT first = 1 if the cards had 100 values initially in the session and is 0 otherwise.

Note the substantial improvement in the R^2 when additional explanatory variables are included, showing the importance of these variables. The regressions clearly indicate that zero bids are substantially less likely when the cards can take on 100 values. Males are significantly more likely to bid zero than are females, and this is true even when we correct for differences in updating skills.[22] The rate of zero bids increases significantly over time in both the 2VT case

---

[22] Overall, of the 130 people for whom we received data on updating questions, 44 answered Question 1 correctly, 47 answered Question 2 correctly, 8 answered Question 3 correctly, and 16 answered Question 4 correctly. 60 people answered no questions correctly, 37 answered one question correctly, 23 answered two questions correctly, 8

and the 100VT whenever the participant first experiences a limited-values treatment. However, we note that when the 100VT is first in a session, the net time trend for zero bids is actually substantially *negative*. It does appear that there is learning over time in the 100VT, but only through the channel of experience with fewer values. Finally, the number of questions answered correctly is indeed a strong predictor of zero-bid rates.

When we instead consider the actual bid made, we see that these are significantly higher in the 2VT than the 100VT, so that people are actually losing more money in this case. This is a nice demonstration that a little sophistication can hurt: A very naïve bidder, who bids, say, 68 in the 100VT, loses 17 on average, while a fully-rational bidder breaks even with a zero bid. A somewhat sophisticated bidder who recognizes that bidding 68 makes no sense but cannot see that 0 is better than 99, loses 24.75! We see little change in the average bid over time, except that there are negative time trends in the 2VT and the 100VT when it is the 2nd treatment in a session. Males' average bids are significantly lower than females' average bids in all specifications, and answering more questions correctly translates into lower average bids. It is interesting to note that the gender effect persists even where we account for the number of questions answered correctly, in specifications (3) and (6).

Finally, we examine the dynamics of bidding behavior with respect to the payoff outcome in the previous period. Do people update their bids based on their previous payoff success, as would seem consistent with most learning models? Table 3 shows how bidding behavior responds to positive, negative or zero lagged outcomes:

---

answered three questions correctly, and 2 answered all four questions correctly. On average, males answered 0.98 questions correctly, while females answered 0.82 questions correctly.

Table 3: Outcomes and Bid Changes

	# times bid raised	# times bid unchanged	# times bid lowered
Positive outcomes			
100VT	201 (28.6%)	267 (38.0%)	235 (33.4%)
4VT	61 (24.9%)	128 (52.2%)	56 (22.9%)
2 VT	53 (5.6%)	781 (76.7%)	235 (17.7%)
Overall	315 (15.6%)	1176 (58.3%)	526 (26.1%)
Negative outcomes			
100 VT	307 (20.2%)	389 (25.6%)	825 (54.2%)
4VT	101(22.7%)	186 (41.9%)	157 (35.4%)
2 VT	220 (15.1%)	850 (58.4%)	386 (26.5%)
Overall	628 (18.4%)	1425 (41.7%)	1368 (40.0%)
Zero outcomes			
100 VT	1200 (31.4%)	2042 (53.4%)	582 (15.2%)
4VT	189 (19.2%)	664 (67.5%)	130 (13.2%)
2 VT	427 (24.7%)	1358 (71.5%)	66 (3.8%)
Overall	1816 (27.3%)	4064 (61.0%)	778 (11.7%)

There are some clear patterns that emerge. After a negative outcome in the 100VT, it is not surprising that a bidder is much more likely to lower the bid than to raise it; however, even after a positive outcome in the 100VT, people are slightly more likely to lower bids than to raise them. A similar pattern holds in the 4VT, although in this case bids are raised slightly more often than they are lowered after positive outcomes. On the other hand, bidders are much more likely to raise the bid than to lower it after a zero outcome, where the bid was insufficient.[23] People are substantially more likely to leave the bid unchanged in the 4VT and the 2VT.[24]

---

[23] There are also a few instances where 1.5 times the value of the card was exactly equal to the (positive) bid made.
[24] While the 2VT data also have these features, they are difficult to interpret, for two reasons: 1) Positive and negative outcomes occur only with positive bids, which are usually 99 and not likely to be raised. 2) Zero outcomes occur only with zero bids, and these cannot be lowered.

There appears to be a tension between a tendency to lower bids after a loss and raise bids after the previous bid was less than the revealed value.

We can also pursue a more formal analysis, using regressions for how the bid made depends on the previous outcome. Table 4 presents two random-effects specifications:

Table 4: Bidding Behavior and Lagged Outcomes

Independent variables	Dependent variable	
	(1) Bid change	(2) Bid change
Constant	3.73* (1.91)	4.39*** (1.97)
2VT	8.41*** (1.96)	8.59*** (1.95)
100 VT	-0.873 (1.97)	-0.325 (1.95)
Male	-2.83** (1.22)	-10.53** (1.93)
Lagged outcome	0.555*** (.009)	0.605*** (.010)
Male*Lagged Outcome	-	-0.443*** (.028)
Zero outcome	5.16*** (1.25)	5.05*** (1.58)
Male*Zero Outcome	-	6.43*** (2.51)
Period	-0.064* (.038)	-0.075** (.037)
N	12095	12095
R^2	.246	.262

Standard errors are in parentheses; *** indicates significance at $p = 0.01$, ** indicates significance at $p = 0.05$, and * indicates significance at $p = 0.10$ (all two-tailed tests). 100 VT = 1 if the cards had 100 values and is 0 otherwise.

One's bid is seen to depend heavily on the payoff outcome from the previous period; the worse the previous outcome, the more the current bid is reduced. In addition, bidders do tend to raise their bids significantly when these were 'unsuccessful' in the previous period. Finally,

while the lagged outcome is significant for both females and males, specification (2) shows that females are substantially more responsive to the lagged outcome than males, so that perhaps there is more learning for females.

5. Conclusion

We find that overbidding persists even when we cast the take-over game as an individual-choice task. We see that the 80% of all bids are positive when there are 100 possible values, 69% are positive when there are four possible values, and 63% are positive when there are only two values. Few people appear to be able to grasp the notion that it is necessary to update their priors about the relevant distribution, although there is more comprehension in the simpler environments.

Indeed, avoiding the WC in the 2VT involves two steps: First, one should realize that bids other than 0 or 99 are nonsensical; and the great majority of participants understand this. Second, one can readily calculate that the expected return from bidding 99 is negative, as there is a 50% chance of losing 99 and a 50% chance of gaining 49.5. Here is where people seem to be unable to see the light, although it dawns on some. The problem is similar with the 4VT, although ordinal logic can only be applied for bids between each pair of possible values. It is particularly interesting that exposure to the two-values or four-values case leads to learning over time in a subsequent 100VT, even though there is no improvement over time in bidding when the 100VT comes first in a session.

Regarding the issue of increased variance vs. decreased complexity, one can interpret the evidence from the 2VT as either exacerbating the WC (greater losses) or reducing it (higher zero-bid rates). We feel that the latter is the more natural interpretation, particularly given the

benchmark of higher expected losses for random bidders in the 2VT, and suspect that many participants may be just a small step away from figuring out that the expected payoff from bidding 99 is negative. In the 4VT, losses are smaller than in the 2VT, even though the zero-bid rate is slightly lower. Nevertheless, this rate is not dramatically lower than in the 2VT, so that the substantial proportion of zero bids in the 2VT cannot be seen to be purely an artifact of there being only two 'reasonable' values (even though this notion of 'reasonable' requires some sophistication from the participants).

Finally, we observe a strong gender effect, with females making higher average bids and a lower proportion of zero bids in all cases. Furthermore, this gender difference persists even when we account for statistical sophistication by including questions on Bayesian updating. Nevertheless, females are more sensitive to lagged outcomes than males are, so that this gender effect may diminish over a substantial period of time.

References

Akerlof, G. (1970), "The Market for Lemons: Qualitative Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, **89**, 488-500.

Ball, S. B., M.H. Bazerman, and J.S. Carroll (1991), "An Evaluation of Learning in the Bilateral Winner's Curse," *Organizational Behavior and Human Decision Processes*, **48**, 1-22.

Bazerman, M.H. and W.F. Samuelson (1983), "I Won the Auction But Don't Want the Prize," *Journal of Conflict Resolution*, **27**, 618-634.

Bereby-Meyer, Yoella and Brit Grosskopf, (2002), "Overcoming the Winner's Curse: An Adaptive Learning Perspective" WP, Harvard Business School.

Bereby-Meyer, Y., B. Grosskopf, and M. Bazerman (2003), "On the Robustness of the Winner's Curse Phenomenon," WP, Harvard Business School.

Blecherman, B. and C.F. Camerer (1998), "Is There a Winner's Curse in the Market for Baseball Players?" mimeograph, Brooklyn Polytechnic University, Brooklyn, NY.

Campbell, Colin, John Kagel, and Dan Levin (1999), "The Winner's Curse and Public Information in Common Value Auctions: Reply," *American Economic Review*, **89**, 325-334.

Capen, E.C., R.V. Clapp, and W.M. Campbell (1971), "Competitive Bidding in High-Risk Situations," *Journal of Petroleum Technology*, **23**, 641-653.

Carroll, J., M. Bazerman, and R. Maury (1988), "Negotiator Cognitions: A Descriptive Approach to Negotiators' Understanding of their Opponents," *Organizational Behavior and Human Decision Processes*, **41**, 351-370.

Cassing, James and Richard W. Douglas (1980), "Implications of the Auction Mechanism in Baseball's Free agent Draft," *Southern Economic Review*, **47**, 110-121.

Charness, Gary and Dan Levin (forthcoming), "When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Psychological Affect", *American Economic Review*.

Cox, James, Sam Dinkin, and James T. Swarthout (2001), "Endogenous Entry and Exit in Common Value Auctions," *Experimental Economics*, **4**, 163-181.

Dessauer, John P., (1981) *Book Publishing*, New York: Bowker.

Dyer, D., Kagel, J. H., and D. Levin. 1989. "A Comparison of Naive and Experienced Bidders in Common Value Offer Auctions: A Laboratory Analysis," *Economic Journal*, 99:108-15.

Eyster, E. and M. Rabin (2002), "Cursed Equilibrium," WP, E02-320, University of California at Berkeley, Berkeley, CA.

Feddersen, T. and W. Pesendorfer (1998), "Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts Under Strategic Voting," *American Political Science Review*, **92**, 23-36.

_____ and _____ (1999), "Elections, Information Aggregation, and Strategic Voting," Proceedings of the National Academy of Science, **96**, 10572-10574.

Gneezy, U. and J. Potters (1997), "An Experiment on Risk Taking and Evaluation Periods," *Quarterly Journal of Economics*, **112**, 631-645.

Grosskopf, Brit, Bereby-Meyer Yoella and Bazerman Max, (2003), "On the Robustness of the Winner's Curse Phenomenon," WP, Harvard Business School.

Hendricks, K., R. Porter, and B. Boudreau (1987), "Information and Returns in OCS Auctions, 1954-1969," *Journal of Industrial Economics*, **35**, 517-42.

Holt, C.A. Jr., and R. Sherman (1994), "The Loser's Curse and Bidder's Bias," *American Economic Review*, **84**, 642-652.

Kagel, J. H. and D. Levin (1986), "The Winner's Curse and Public Information in Common Value Auctions," *American Economic Review*, **76**, 894-920.

_____ and _____. (1991), "The Winner's Curse and Public Information in Common Value Auctions: Reply," *American Economic Review*, **81**, 362-69.

_____and_____ (1999), "Common Value Auctions with Insider Information," *Econometrica*, **67**, 1219-1238.

_____ and _____ (2002), *Common Value Auctions and the Winner's Curse*, Princeton: Princeton University Press.

_____ , _____ and R. Harstad (1995), "Comparative Static Effects of Number of Bidders and Public Information on Behavior in Second-Price Common Value Auctions," *International Journal of Game Theory,* **24**, 293-319.

Levin D., Kagel J. H. and J. F. Richard (1996), "Revenue Effects and Information Processing in English Common Value Auctions" *American Economic Review*, **86**, 442-460.

Lorenz, J., and E.L. Dougherty (1983), "Bonus Bidding and Bottom Lines: Federal Off-shore Oil and Gas," SPE 12024, 58th Annual Fall Technical Conference.

Leitzinger, Jeffrey J. and Joseph E. Stiglitz (1984), "Information Externalities in Oil and Gas Leasing," *Contemporary Policy Issues*, **5**, 44-57.

Levis, Mario (1990), "The Winner's Curse Problem, Interest Costs and the Underpricing of Initial Public Offerings," *Economic Journal*, **100**, 76-89.

Mead, W. J., A. Moseidjord, and P.E. Sorensen (1983), "The Rate of Return Earned by Leases Under Cash Bonus Bidding in OCS Oil and Gas Leases," *Energy Journal*, **4**, 37-52.

_____, _____, and _____ (1984), "Competitive Bidding Under Asymmetrical Information: Behavior and Performance in Gulf of Mexico Drainage Lease Sales, 1954-1969," *Review of Economics and Statistics*, **66**, 505-08.

Rock, Kevin (1986), "Why New Issues are Underpriced," *Journal of Financial Economics*, **15**, 187-212.

Roll, Richard (1986), "The Hubris Hypothesis of Corporate Takeovers," *Journal of Business*, **59**, 197-216.

Samuelson, W.F. and M.H. Bazerman (1985), "The Winner's Curse in Bilateral Negotiations," in V.L. Smith (ed.), *Research in Experimental Economics*, Vol. 3, Greenwich, CT: JAI Press.

Selten, Reinhard, Klaus Abbink, and Ricarda Cox (2001), "Learning Direction Theory and the Winner's Curse," Bonn Econ Discussion Papers bgse10_2001, University of Bonn.

Tor, Avishalom and Max H. Bazerman, "Focusing Failures in Competitive Environments: Explaining Decision Errors in the Monty Hall Game, The Acquiring a Company Problem and Multi-party Ultimatums," WP, Harvard Business School.

Appendix A - Instructions

Welcome to our experiment. You will receive $5 for showing up, regardless of the results. In addition, you receive a $15 endowment, with which you make bids on *cards* with some distribution of values.

You will see 100 cards on your screen, in a 10x10 array. They will have values that range between 0 and 99, inclusive. You will see all of these values before beginning, and will then click a button to flip the cards over and start the game. The cards are re-shuffled after every action.

Each period you will choose a bid (a non-negative integer) and will then select one of the cards; its value will then be shown.

If your bid is greater than or equal to the card's value, you receive 150% of this card's value. If your bid is less than the card's value, nothing happens; that is, you earn zero in such rounds.

> If, for example, you bid 42:

1. Suppose the value of your selected card is 36, then its worth to you is $36 \times 1.5 = 54$; since your bid was 42, your profit is 12.
2. Suppose the value of your selected card is 20, then its worth to you is $20 \times 1.5 = 30$; since your bid was 42 your profit is -12.
3. Suppose the value of your selected card is 67, then no transaction takes place and your profit is 0.

Your results will be displayed on the right side of the screen after you choose the card.

There will be a total of 60 periods in the session, consisting of two segments of 30 periods. The distribution of the card's values will change after 30 periods, and the new values will be shown to you. You then click the button to flip them over and begin the 2^{nd} segment.

A history of the results of all previous periods (showing period, bid, value, profit) will be displayed at the bottom of the screen.

At the end of the 60 periods, you will enter a name and then submit your results.

We will randomly choose one period from each 10-period block for payment, so that only six periods will actually count towards monetary payoffs. These periods will be chosen randomly at the end of the session, and will be shown on your screen. We will then net out your gains and/or losses from these six periods against your original endowment, converting them to actual dollars at the rate of $0.10 for each unit.

At the end of the experiment, we will pay each participant individually and privately.

Thank you again for your participation in our research.

Appendix B – Questions on Bayesian Updating

Consider two machines placed in two sides of large production hall, left side = L and right side = R.

The two machines produce rings, good ones and bad ones. Each ring that comes from the left machine, L, has a 50% chance of being a good ring and a 50% chance of being a bad ring. Each ring that comes from the right machine, R, has a 75% chance of being a good ring and a 25% chance of being a bad ring.

In each of the following 4 questions you will observe some ring(s) that are the output of one of the two machines. After the information is given, you are asked to mark one (and only one) of the probabilities that you think is closest to the true probability that the L = left machine produced the ring(s), given your observations.

Each correct answer pays $0.50.

Q1. You observe one **good ring**. What is your best assessment of the probability that the left machine produced this ring?

1.0	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.2	0.1	0.0

Q2. You observe one **bad ring**. What is your best assessment of the probability that the left machine produced this ring?

1.0	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.2	0.1	0.0

Q3. You observe six rings, and **all six are bad.** What is your best assessment of the probability that the left machine produced these rings?

1.0	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.2	0.1	0.0

Q4. You observe six rings, of which **three are good** and **three are bad**. What is your best assessment of the probability that the left machine produced these rings?

1.0	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.2	0.1	0.0