

The value of victory: social origins of the winner's curse in common value auctions

Wouter van den Bos^{*1,2}, Jian Li³, Tatiana Lau¹, Eric Maskin⁴,
Jonathan D. Cohen^{1,5}, P. Read Montague³, and Samuel M. McClure⁶

¹ Department of Psychology and
Center for the Study of Brain, Mind, & Behavior, Princeton University

² Institute for Psychological Research, Leiden University

³ Department of Neuroscience and Human Neuroimaging Laboratory,
Baylor College of Medicine

⁴ Institute for Advanced Study, Princeton

⁵ Department of Psychiatry, University of Pittsburgh

⁶ Department of Psychology, Stanford University

Abstract

Auctions, normally considered as devices facilitating trade, also provide a way to probe mechanisms governing one's valuation of some good or action. One of the most intriguing phenomena in auction behavior is the winner's curse — the strong tendency of participants to bid more than rational agent theory prescribes, often at a significant loss. The prevailing explanation suggests that humans have limited cognitive abilities that make estimating the correct bid difficult, if not impossible. Using a series of auction structures, we found that bidding approaches rational agent predictions when participants compete against a computer. However, the winner's curse appears when participants compete against other humans, even when cognitive demands for the correct bidding strategy are removed. These results suggest the humans assign significant future value to victories over human but not over computer opponents even though such victories may incur immediate losses, and that this valuation anomaly is the origin of apparently irrational behavior.

Keywords: auctions, winner's curse, bounded rationality.

1 Introduction

Auctions of various types constitute a ubiquitous mechanism for buying and selling items of limited quantity or availability. Given the commonality of auctions, it is noteworthy that under fairly general conditions, even highly experienced bidders tend to suffer net losses (Kagel & Richard, 2001). This fact has come to be called the winner's curse and was first identified in auctions for drilling rights in the Gulf of Mexico (Capen, Clapp, & Campbell, 1971). The phenomenon has since been reported in a range of other field settings (Ashenfelter & Genesove, 1992; Blecherman & Camerer, 1998; Cassing & Douglas, 1980; Dessauer, 1981) as well as laboratory studies (Bazerman & Samuelson, 1983; Kagel & Levin,

1986; Kagel, Levin, Battalio, & Meyer, 1986). Experiments have shown that naive bidders initially incur large losses that decline over time, but the “curse” nonetheless persists even for very experienced (Garvin & Kagel, 1994; Milgrom & Weber, 1982) or professional (Dyer, Kagel, & Levin, 1989) auction participants.

The winner's curse arises in auctions for items of fixed, but unknown, value (known as common value auctions). Oil drilling rights satisfy these conditions because the amount of oil in a region (and hence its market value) is the same for all bidders, yet cannot be precisely estimated. To understand the source of the winner's curse, begin by assuming that estimates of value for each bidder are correct on average with some amount of variance. If participants bid their estimated value, the winner of the auction will be that participant with the most optimistic estimate. Since estimates are distributed around the true value, the most optimistic estimate will generally be an

^{*}This work was supported by NIDA grant R01 DA11723 (PRM) and NIA grant R01 AG030310 (JDC). Correspondence to: WVDB (wbos@fsw.leidenuniv.nl), SMM (smcclure@stanford.edu).

overestimate of the true value and a net loss will therefore occur for that bidder. Hence, on average the winner is cursed by the statistical fact that their estimate is more likely than not to be greater than the true value.

To avoid the winner's curse it is necessary to modify one's bid beyond one's estimate of the true value to account for the conditional probability of winning the auction. That is, a "good" bid ought to be sufficiently less than one's estimate in order to acknowledge the fact that winning is most likely to occur for an overestimate. In practice, auction participants do modify their bids strategically, but the correction is not sufficient to avoid the winner's curse (Eyster & Rabin, 2005; Kagel & Richard, 2001; Kagel & Levin, 2002).

There have been several proposed mechanisms for the winner's curse (e.g., Eyster & Rabin, 2005; Parlour, Prashnikar, & Rajan, 2007). These explanations generally propose that bidders fail to perform rationally due to cognitive limitations (Fudenberg, 2006). The hypothesis is that people understand they must submit bids less than their estimates, but they are unable to accurately calculate exactly how much less to bid. In this study, we examine this hypothesis and find no support for it. We first present data demonstrating that the winner's curse is not a consequence of limited cognitive abilities. We then demonstrate that the curse depends instead on the social nature of the auction environment.

Previous work on the psychology of auctions has demonstrated that social influences have a significant effect on bidding. In escalating auctions, in which participants bid sequentially until a single high bidder remains, people are subject to intense emotions that impede rationalization and may lead to extremely high bids (Ku et al., 2005). This phenomenon, known as "auction fever," is likely to be related to the winner's curse, according to our thesis. The winner's curse has previously been thought to be an entirely separate phenomenon because the uncertainty inherent in common value auctions presumably make cognitive demands overwhelming (Ku et al., 2005; Kagel and Levin, 2002). Additionally, in our experiments the social environment is minimal. Each participant submits a sealed bid essentially in isolation and the highest submitted bid wins, leaving no opportunity for competitive fire to escalate with the progression of the auction. Our findings address both of these points. First, despite the challenges posed by lack of information in common value auctions, people are able to converge to stable bidding strategies in less than 50 trials that remain constant when cognitive demands are eliminated. Second, we find that the winner's curse is in fact strongly dependent on social context. We conclude that competitive arousal, and not cognitive limitations, underlies the winner's curse.

2 Method

2.1 Participants

The studies were conducted at Princeton University in Princeton, NJ and Baylor College of Medicine in Houston, TX, where respectively 47 and 48 volunteers participated in the experiments. The average age of the group was 26.39 years (S.E. 0.90), and consisted of 33 male and 62 female participants. Although volunteers were recruited from separate participant pools they were instructed identically, but separately, at each institution (instructors presented an identical set of PowerPoint slides as instructions for the task). All participants passed a mathematics quiz given after the experiment to ensure that all had the quantitative skills necessary for this experiment.

2.2 Procedure

The experiments were run in three conditions with separate groups of participants. In all conditions, participants first engaged in auctions of 5 or 6 participants in a baseline condition that we call Human/Naive (Experiment 1). The behavior of the three participant groups of in this initial experiment was undistinguishable and are reported in aggregate.

After this first experiment, participants were excused and recalled to the laboratory after a 2 week period. At this point, all participants were given instructions about the winner's curse and instructed how to calculate the risk-neutral Nash equilibrium (RNNE) bidding strategy that maximizes expected profits (see below). All participants were given written tests to confirm that they understood the RNNE strategy and could compute RNNE bids easily. It was after this initial auction experience and instruction that the three conditions differed in experimental procedure.

In the first of the three conditions, which we call Human/Expert (Experiment 2; $n=28$), participants engaged in another set of auctions against 5 or 6 other people. For the other two conditions, participants bid against computer algorithms whose bidding strategy was explained to the participants. For the second condition, called Computer/Emulation (Experiment 3; $n=38$), the computer opponents bid by drawing from the distribution of bids submitted by human participants in Experiment 2. In the final condition, Computer/RNNE (Experiment 4; $n=28$), computer opponents bid the RNNE strategy explained below.

The procedures during the auctions was identical in all experiments (see Appendix A for task instructions). Each experiment consisted of fifty consecutive sealed bid auctions. In each auction round, participants were provided

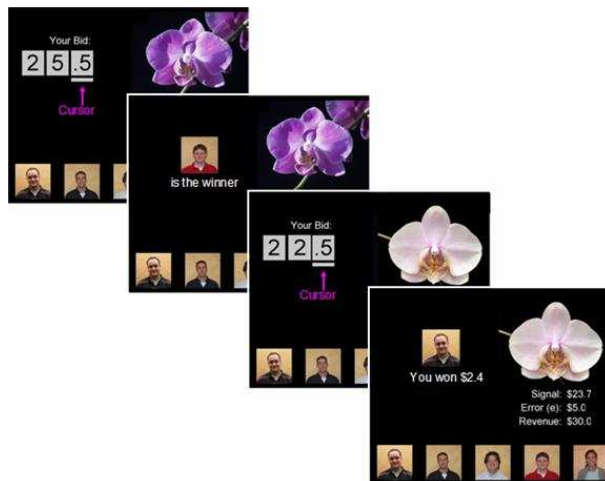


Figure 1: For each auction, participants were told their personal estimate of the item's value (x_i), the error term ε , and their current revenue. Pictures of the other participants were displayed on the bottom of the screen. For experiments 3 and 4, the participant photos were replaced with icons of computers. In each auction, all participants simultaneously entered their bids; individual bids were never revealed to other participants. After all bids were submitted, the winning bidder was revealed to all with no information given about the amount of money won or lost in the auction. In each of the four experiments, participants completed 50 rounds of auctions with random estimates and errors determined on each round.

with two pieces of information about the value of the item under auction (see Figure 1). First, they were provided independent estimates of the value of the item under auction, which we call x_i in the remainder of the paper. They were also instructed that estimates were drawn from a uniform distribution with maximum error ε around the true common value x_0 . Before the start of the experiment, participants were additionally instructed that x_0 was randomly drawn from uniform distribution with maximum and minimum values x_H and x_L , respectively.

All participants were endowed with \$30 at the beginning of the experiment. The winner of each auction round earned $x_0 - b$, where b is the size of the winning bid. This amount (which could be a gain or a loss) was added to his or her revenue. All other participants earned \$0. Fifteen of the participants lost all of their endowment during the course of the Human/Naive experiment. We allowed participants to continue bidding in this case to preserve the number of participants in the auctions and to ensure equal experience prior to receiving instruction on how to perform in accord with RNNE. Allowing participants to continue bidding after bankruptcy has been found to have no effect on bidding (Lind & Plott, 1991).

In all experiments, the winner of the auction was shown how much they earned or lost, but all other participants were only shown the identity of the winner. This corresponds to Armantier's (2004) minimal information condition, but deviates from other common value auction experiments. In particular, we did not inform participants how much the other participants bid. We also did not show losing bidders the true value of the item, x_0 (cf. Kagel & Levin, 1986).

2.3 RNNE bid strategy

With the information given to participants in the experiments, the risk-neutral Nash equilibrium (RNNE) bidding strategy can be determined. RNNE bidding gives the optimal bidding strategy in the sense that it maximizes expected earnings for all participants. The solution is given by

$$b_{RNNE} = x_i - \varepsilon + Y \tag{1}$$

where

$$Y = \frac{2\varepsilon}{n + 1} \exp\left(\frac{n}{2\varepsilon} [x_i - (x_L - \varepsilon)]\right) \tag{2}$$

and n is the number of bidders (Milgrom & Weber, 1982). This function assumes that participants are risk neutral. When estimates are farther than ε from the bounds on x_0 (which was always true in our experiments), then Y is very close to zero, and can be ignored (Garvin & Kagel, 1994).

As can be seen in Figure 2, participants almost always bid above the Nash equilibrium bidding strategy in Experiment 1. One possibility for why this occurs is that people adopt a "naive" bidding so that bids match their estimates. A continuum can be generated between this strategy and the optimal strategy by expressing bids according to

$$b_i = x_i - (1 - \kappa)\varepsilon \tag{3}$$

where κ captures the degree to which bids exceed the optimal strategy. The Nash equilibrium and naive bidding strategies fall conveniently at $\kappa = 0$ and $\kappa = 1$, respectively. We call κ the "bid factor" in subsequent discussion.

Importantly, for our experiments the RNNE strategy is achieved by bidding $x_i - \varepsilon$. Since both x_i and ε were presented on each round of auction, optimal RNNE bids could be calculated with a single subtraction. In the "expert" conditions in Experiments 2, 3, and 4 below, we ensured that all participants understood the rationale behind bidding below estimates in this manner and that they were able to compute this quantity easily.

In some cases our experimental data were not normally distributed (as determined by Kolmogorov-Smirnov tests). We handle this by reporting the results using both

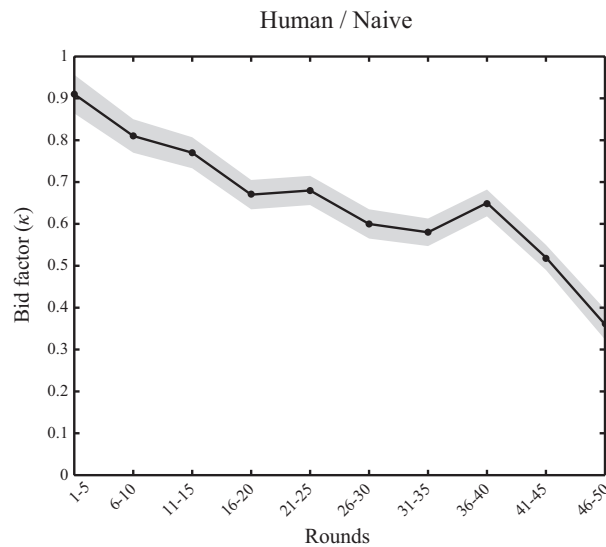


Figure 2: Over-bidding relative to the rational bidding strategy (bid factor) in plotted, averaged over sequential 5 rounds of auctions in the Naive experiment (Experiment 1; gray bars indicate S.E.). An ANOVA with number of rounds as a between participants factor confirmed that the naïve participants learned to decrease the size of their bids ($F(9,86) = 16.509, p < .001$). This learning effect was absent in all follow up experiments ($p > .1$ for Experiments 2, 3, and 4).

non-parametric tests as well as using t tests on data transformed according to

$$\kappa' = \frac{\log\left(\frac{b_i}{b_{RNNE}}\right)}{\log\left(\frac{x_i}{b_{RNNE}}\right)} \quad (4)$$

This allows us to take advantage of the power of the parametric tests while also avoiding problems based on inappropriate assumptions about the structure of the data. As seen in the results, the exact statistical test used had no effect on our conclusions.

3 Results

3.1 Experiment 1: The winner's curse

The results for Experiment 1 demonstrate that naïve participants consistently bid above the Nash equilibrium (median $\kappa = 0.655$; Figure 2 and Figure 3A) and consistently lose money (Figure 4). Losses are particularly strong in early rounds and diminish with time as participants gain experience (Figure 2). However, the winner's curse (indicated by $\kappa > 0$) persists through all 50 rounds of auctions of our experiment. This is consistent with

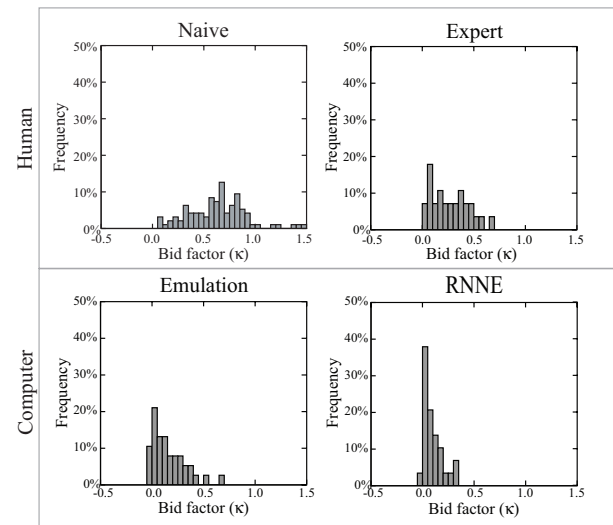


Figure 3: Histograms of the frequency of bid factors for all bid submitted in each of the four experiments. Bids were significantly reduced when participants bid against computer opponents (lower plots). Most participants in the expert auctions (all but top-left plot) appear to have treated the RNNE bidding strategy ($\kappa=0$) as a lower bound for submitted bids. The distribution of bids was positively skewed in these auctions as a consequence.

prior studies demonstrating that the winner's curse is evident even for participants with a large amount of experience (Garvin & Kagel, 1994; Milgrom & Weber, 1982; Dyer et al., 1989).

The size of the winner's curse (given by the magnitude of bids relative to RNNE, i.e. $b - b_{RNNE}$) correlated strongly with the possible error in value estimates (ϵ ; $r = 0.35$; $p < 10^{-10}$). This implies that participants use the error information to scale their as suggested by Equation 3. This finding justifies the use of the bid factor to summarize bidding.

3.2 Experiment 2: Persistence of winner's curse in absence of cognitive demands

For Experiment 2, auctions were composed of a subset of the participants from Experiment 1 who had received instruction about how to maximize expected earnings in the task using the RNNE strategy. Despite being fully competent in implementing this strategy, the winner's curse persisted (median $\kappa = 0.25$), remaining at the same level at which bidding ended in Experiment 1 ($t(1, 27) = 1.383$; $p = 0.178$; comparison with median κ in final 5 rounds in Experiment 1). Interestingly, this bid factor produces an average of no net change in revenue during the course of the experiment ($t(1, 27) = -.238$, $p = 0.813$). With our number of participants, the RNNE

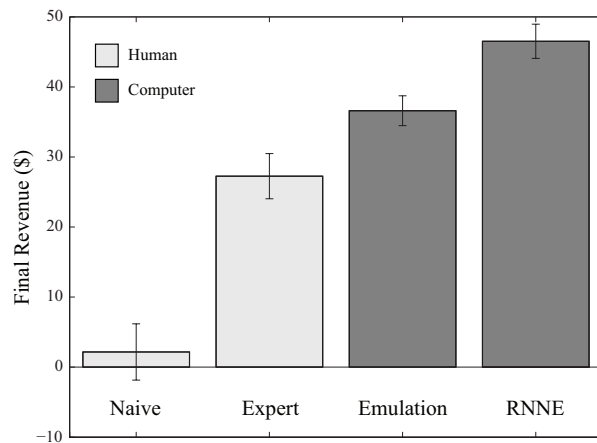


Figure 4: Participants were endowed with \$30 at the beginning of each experiment. The average revenue at the end of 50 rounds of auctions is shown for each of the experiments (error bars represent S.E.).

strategy guarantees no loss of revenue in each auction and expected profits overall. Since participants were instructed in the RNNE bidding strategy and tested to ensure their ability to implement this strategy, the fact that bids were unaffected by this manipulation strongly suggests that factors other than purely cognitive limitations underlie the winner's curse. Instead, post-experiment debriefings strongly suggested that overbidding was driven by the social nature of the auctions. Participants reported that their bidding was driven by a desire to "win" more often than the rationally optimal strategy allowed.

3.3 Experiment 3: Social factor underlie the winner's curse

The final two experiment conditions were designed to test our hypothesis that social context modulates the winner's curse. In both conditions we had participants bid against computer opponents while preserving all other aspects of the task from Experiment 2.

The RNNE strategy assumes that all participants bid optimally. However, in Experiment 2, participant bid with positive bid factor. There are strategic consequences when bidding against overly aggressive opponents. In particular, the RNNE strategy increases when opponents bid above the RNNE solution in Equation 1. For Experiment 3 (Computer/Emulation) we had the computer bid so as to match the distribution of bids submitted by experienced participants in Experiment 2. The strategic effects of the behavior of the other participants is therefore equivalent in this experiment and in Experiment 2.

A different subset of participants from Experiment 1 participated in Experiment 3. Despite the fact that per-

formance by the other participants was preserved in this experiment relative to the Human/Expert condition (Experiment 2), the winner's curse was significantly reduced with the change to computer opponents (Figure 3; median $\kappa = 0.102$, Mann-Whitney U-test, $z = -2.674$, $p = 0.003$, one-sided; $\kappa' : t(1, 64) = 8.527$; $p < 0.001$; two-tailed; Figure 3). Additionally, for the first time participants earned net profits during the experiment ($t(1, 37) = -5.407$, $p < .001$; Figure 4). This finding indicates that social factors play an important role in generating the winner's curse evident in the behavior of the "expert" bidders in Experiment 2.

3.4 Experiment 4: Strategic effects of bidding against participants

The remaining participants from Experiment 1 participated in Experiment 4. In our final experiment, Experiment 4 (Computer/RNNE condition), we had computer opponents bid according to the RNNE strategy (i.e. $x_i - \epsilon$). All other aspects of the task were identical to the Computer/Emulation experiment.

With this final manipulation, bidding was reduced by a small but significant amount compared with Experiment 3 (median $\kappa = 0.07$, Mann-Whitney U-test, $z = -1.778$, $p = 0.04$, one-sided; $\kappa' : t(1, 65) = 2.168$; $p < 0.017$, one-tailed; Figure 3). The fact that bid factors are still above zero seems to be due to the fact that participants' used RNNE bidding as a lower bound and occasionally bid above this value (Figure 3, lower right panel). The distribution of bids is significantly skewed and has a positive median bid factor as a consequence. Nonetheless, the distribution of bids submitted in this experiment indicates that with our instructions participants are able to implement the RNNE bid strategy, but may fail to do so as a consequence of social context and subsequent (small) strategic effects.

4 Discussion

Social factors have been discussed as a potential cause of the winner's curse for some time, but have been dismissed on various grounds (e.g., Holt & Sherman, 1994; Goeree et al., 2002; Ku et al., 2005; but see Delgado, Schotter, Ozbay, & Phelps, 2008)). One attraction of the hypothesis is that it preserves optimality analyses by the inclusion of a couple of additions terms in participants' utility functions. In particular, the utility based on earnings for bidder i is given by

$$U_i = \begin{cases} x_0 - b_i & \text{if } b_i = \max_{j \in n} b_j \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

where b_i is the bid, n is the set of participants in the auction, and x_0 is the common value of the item under auction. If winning and losing affect utility independent of the monetary consequences of the auction such that:

$$U_i = \begin{cases} x_0 - b_i + r_{win} & \text{if } b_i = \max_{j \in n} b_j \\ -r_{lose} & \text{otherwise} \end{cases} \quad (6)$$

where r_{win} and r_{lose} depend on social factors, then optimal bids are increased by an amount equal to $r_{win} + r_{lose}$.

One problem with this explanation is that it implies that bids should be increased by a constant amount relative to RNNE. Thus, variables such as the error (ϵ) should not correlate with the size of bids relative to RNNE. However, in our experiments participants bid farther above RNNE with greater error (a similar finding was reported for private value auctions by Goeree et al., 2002). Unless it is assumed that r_{win} and r_{lose} are also proportional to ϵ then this observation is unexplained by the "utility of winning" in Equation 6. We do not have a complete explanation for this, but note that if r_{win} and r_{lose} do depend on ϵ then it allows the probability of winning to be increased independent of ϵ (see Appendix B). It seems reasonable to presume that competition may drive people to alter behavior with the goal of increasing the chances of winning. Thus, it is not unreasonable that the utility of winning depends on factors that alter the probability of winning as well.

Other work has taken a similar approach to ours to directly measure the utility of winning. In a clever experiment by Holt and Sherman (1994), a separate statistical effect that produces a "loser's curse" was used to offset the winner's curse. With the proper parameters, the two effects should perfectly offset leaving only the utility of winning as a bias on bidding. With this approach, Holt and Sherman concluded that the utility of winning is zero. However, these auctions were of a very different structure than standard common value auctions and in fact only had single participants. The finding that there is no utility of winning when a single person bids against a computer is fully consistent with our results.

Finally, the utility of winning has been dismissed as an explanation for the winner's curse due to the fact that the scarcity of information in common value auctions make the statistical argument seem natural and compelling (Kagel & Levin, 2002; Ku et al., 2005). Additionally, compared with other auction structures in which social factors are especially salient, sealed bid auctions do not seem to evoke the same level of competitive exhilaration that is presumed to underlie "auction fever" (Ku et al., 2005). We can speak to this in two ways. First, we note that people seem perfectly capable of adapting to the winner's curse with a moderate amount of experience. In Experiment 1, bidding reached asymptote after

less than 50 trials. Armantier (2004) proposed that people learn through reinforcement learning, which predicts the gradual learning curve in Figure 2. Thus, while calculating the appropriate bid shading is possible using mathematical reasoning, it appears that people may actually do something far easier. In particular, the reinforcement learning approach suggests that people learn from trial-and-error and gradually reduce their bids until they reach a desired level of average returns. The naturalness of the statistical problem that underlies the winner's curse may be entirely unrelated to how people respond to incurred losses and gains.

Second, it is certainly true that the social context in sealed bid common value auctions is not as enveloping as, for example, English outcry auctions in which the escalation of offers is publicly observed. However, our results indicate that the mere presence of human competitors in sufficient to increase bidding. Furthermore, the social context is certainly more powerful in our experiments than other manipulations that are known to affect behavior in economic games. For example, in Miller et al. (1998), simply telling someone that a fictional partner shares one's birthday increases cooperation in a Prisoner's Dilemma game. We had participants meet each other prior to the experiment and displayed pictures of participants' faces during the experiment, which seems like a richer social context than Miller et al.

Our intention in designing the experiments was to maximize the feelings associated with being in a social context. This was the motivation behind displaying the winner's picture after each auction. Other investigators have made efforts to minimize this effect with the goal of equalizing behavior when playing against human and computer competitors (Walker et al., 1987). Exactly what manipulations are necessary to create a competitive social context remains a difficult future challenge. However, it is important to note here that it seems certainly possible that subtle variations on our experimental design may eliminate competitive motivations that we find to underlie the winner's curse. By contrast, our results suggest other conditions under which the winner's curse may become especially pronounced. For example, situations in which auction winners are highly visible, such as occur in sports free agency, may incur especially high utility for winning.

Finally, the ability to measure individual susceptibility to the Winner's Curse offers a mechanism to measure socially-derived value in units of dollars. This, in turn, may enable finer investigation of psychiatric disorders with associated social dysfunctions. Behavioral economics experiments of this sort have recently been used to probe psychiatric disorder in this manner (King-Casas et al., 2008); common value auctions may be another tool in a behavioral economics-derived psychiatry battery.

References

- Armantier, O. (2004). Does observation influence learning? *Games and Economic Behavior*, 46, 221–239.
- Ashenfelter, O., & Genesove, D. (1992). Testing for price anomalies in real-estate auctions. *American Economic Review*, 82, 501–505.
- Bazerman, M. H., & Samuelson, W. F. (1983). I won the auction but don't want the prize. *Journal of Conflict Resolution*, 27, 618–634.
- Blecherman, B., & Camerer, C. F. (1998). Is there a winner's curse in the market for baseball players? Brooklyn Polytechnic University, mimeograph, Brooklyn, NY.
- Capen, E. C., Clapp, R. V., & Campbell, W. M. (1971). Competitive bidding in high risk situations. *Journal of Petroleum Technology*, 23, 641–53.
- Cassing, J., & Douglas, R. W. (1980). Implications of the auction mechanism in baseball's free agent draft. *Southern Economic Journal*, 47, 110–121.
- Delgado, M.R., Schotter, A., Ozbay, E.Y., & Phelps, E.A. (2008). Understanding overbidding: using the neural circuitry of reward to design economic auctions. *Science*, 321, 1849–1852.
- Dessauer, J. P. (1981). *Book publishing: What it is, what it does*. Bowker: New York.
- Dyer, D., Kagel, J. H., & Levin, D. (1989). A comparison of naive and experienced bidders in common value offer auctions: A laboratory analysis. *The Economic Journal*, 99, 108–115. E
- Eyster, E., & Rabin, M. (2005). Cursed equilibrium. *Econometrica*, 73, 1623–1672.
- Fudenberg, D. (2006). Advancing beyond advances in behavioral economics. *Journal of Economic Literature*, 44, 694–711.
- Garvin, S., & Kagel, J. H. (1994). Learning in common value auctions: Some initial observations. *Journal of Economic Behavior & Organization*, 25, 351–372.
- Goeree, J. C., & Offerman, T. (2002). Efficiency in auctions with private and common values: an experimental study. *American Economic Review*, 92, 625–643.
- Holt, C. A., & Sherman, R. (1994). The loser's curse. *American Economic Review*, 84, 642–652.
- Kagel, J. H., & Richard, J.-F. (2001). Super-experienced bidders in first-price common-value auctions: Rules of thumb, Nash equilibrium bidding, and the winner's curse. *The Review of Economics and Statistics*, 83, 408–419.
- Kagel, J. H., & Levin, D. (1986). The winner's curse and public information in common value auctions. *American Economic Review*, 76, 894–920.
- Kagel, J. H., Levin, D., Battalio, R. C., & Meyer, D. J. (1989). First price common value auctions: Bidder behavior and the "winner's curse." *Economic Inquiry*, 27, 241–258.
- Kagel, J. H., & Levin, D. (2002). *Common Value Auctions and the Winner's Curse*. Princeton University Press: Princeton, NJ.
- King-Casas, B., Sharp, C., Lomax-Bream, L., Lohrenz, T. Fonagy, P. & Montague, P.R. (2008) The rupture and repair of cooperation in borderline personality disorder. *Science*, 321, 806–810.
- Ku, G., Malhotra, D., & Murnighan, J. K. (2005). Towards a competitive arousal model of decision-making: A study of auction fever in live and internet auctions. *Organizational Behavior and Human Decision Processes*, 96, 89–103.
- Lind, B., & Plott, C. R. (1991). The winner's curse: experiments with buyers and with sellers. *American Economic Review*, 81, 335–346.
- Milgrom, P. R., & Weber, R. J. (1982). A theory of auctions and competitive bidding. *Econometrica*, 50, 1089–1122.
- Miller, D.T., Downs, J.S., & Prentice, D.A. (1998). Minimal conditions for the creation of a unit relationship: the social bond between birthdaymates. *European Journal of Social Psychology*, 28, 475–481.
- Parlour, C. A., Prasnika, V. & Rajan, U. (2007). Compensating for the winner's curse: Experimental evidence. *Games and Economic Behavior*, 60, 339–356.
- Walker, J. M., Smith, V. L. & Cox, J. C. (1987). Bidding behavior in first price sealed bid auctions: use of computerized Nash competitors. *Economic Letters*, 23, 239–244.

Appendix A: Task instructions

Sealed bid auction experiment

For experiments against human participants:

During this experiment you will participate in multiple Sealed Bid Auctions with 5 other players.

For experiments against computer opponents:

During this experiment you will participate in multiple Sealed Bid Auctions against 5 computer players. These computer players will be playing according to a pre-set method that is explained later in these instructions.

Sealed bid auction

A Sealed Bid Auction (SBA) is an auction where several bidders simultaneously submit bids to the auctioneer without knowledge of the amount bid by other participants. The person who submits the highest bid is the winner of the auction. Subsequently, the winner will pay the price of his bid for the object under auction. Usually,

and also in this experiment, the auction is also a common value auction — meaning that the item up for auction has a fixed, but unknown value.

Common value auction

In all auctions in this experiment, the goods will be worth exactly the same for each of the participants. Usually the objects that are bought at a common value auction are sold at a later stage. The true value of the objects is the future resale value of the objects, and is the same for everybody. However, it is important to understand that this true value *will not be known by any of the participants*. Instead, every participant has a privately known *estimate of the true value*. That private estimation is called **the signal**, and it is given by the experimenter. Naturally the signal can be *higher or lower* than the true value.

The experiment

In this experiment you will be participating in a series of sealed bid auctions that will be auctioning different types of flowers. In each auction a new type of flower will be presented. Each flower will be shown with two additional pieces of information:

- (1) your PERSONAL SIGNAL of the value of the flower,
- (2) an ERROR TERM that indicates how far your signal can be from the *true value* of the flower.

Based on this information you will know only imprecisely the actual, true value of the flower being auctioned. More detail on the signal and error information is below.

Each player will begin with the same amount of money. When the new object (a flower) is presented on your computer screen, you can make your bid; adjust the digits so that they represent the bid you want to make and then submit it. When all sealed bids are received the winner is determined. The revenue (which is explained later) of that bid (which can be negative or positive) will be added to the winner's revenue, and the name and picture of the winner will be displayed on every screen. The next round will start after a few seconds.

Revenue

The earnings on each auction are determined as the difference between the amount paid (winning bid) and the true value of the flower. Since the true value is not known until the end of the auction, it is to your advantage to make bids that you believe to be *less* than the (unknown) true value of the flowers. As long as you bid less than the true value of the items, you will make profits on the auctions. If you win an auction with a bid that is above the true value of the flower, you will lose certain amount of money that is

equal to the difference between your winning bid and the true value of the flower.

The Signal (s) and the Error term (e)

Each participant will receive a private signal (s) that is based upon the real true value of the flowers for sale (the private signal will be different for each participant). For each round of the auction, an error range (e) is picked. The error term (e) is displayed for each round and indicates how far estimates can be from the true value of the flower being auctioned. All players receive the same error range (e) for any specific round. The true value of the auctioned flower will be within the error range of the private signal (s±e).

Example:

True Value	\$16	Same for everybody, but not known
Error (e)	\$2	Indicates that individual signals of the true value can be between \$14 and \$18

A random personal error is picked for each round. Unlike the error range (e), this personal error is different for each player and will remain unknown.

For instance, assume that the error picked for you is -\$1.25. That means that your signal will be:

$$\$16 - \$1.25 = \$14.75$$

Keep in mind that you do not know your personal error. However, you do know that the real common value must be somewhere in between your signal (\$14.75) minus maximal error (\$14.75-\$2=\$12.75) and your signal plus maximal error (\$14.75+\$2=\$16.75), thus in the range of \$12.75 to \$16.75.

The distribution of the error term is uniform; this means that real value of the auctioned flower is equally possible between s±e (signal ± error range). A larger error is as likely as a small error.

For computer/emulation experiment the following paragraph was included here:

Computer players will also play under these rules. Each one will be assigned different personal errors and signals. Each computer player, however, will submit bids based on how previous experiment participants performed with your same level of experience.

For computer/rnne experiment the following paragraph was included here:

Computer players will also play under these rules. Each one will be assigned different personal errors and signals.

Each computer player, however, will automatically submit the Nash equilibrium bid (see below) every round.

After each round of bidding, the picture and name of the winner for that specific round will be revealed to all participants, though none of the participants knows exactly how much revenue the winner earned for that round.

The range of the common value

There is also a minimum possible real value and a maximum real value for the flowers sold on the auction. All goods in this experiment will have a real value between \$10 and \$48.

What to do?

- To change the digit bid amount: Press 1
- To move the cursor to another digit: Press 2
- To enter a bid: Press 3 or 4

Please use your left hand for buttons 1 and 2 on the keyboard and you right hand for button 3 or 4 on the numpad (check numlock is on). This will prevent to submit bids by mistake.

Every player must try to make as much revenue as possible during the course of the experiment. At the end there will be an extra monetary reward based upon how much revenue you made.

- Everybody will earn at least \$15 dollars for playing the game.
- If you end up with less virtual money than you start with in the game you will not make any extra money.
- If you end up with negative revenue (below zero!), money will be subtracted from your original \$15 dollars.
- If you make money in the game you can make up to \$30 dollars total.

Only at the end of the experiment will the total revenue of each player will be revealed to everyone else.

Remember that if you never win an auction, you will not garner any additional revenue. However, bidding too high and paying more than the real common value means losing revenue!

At the end you will be asked how well you believe you did relative to the other players, so try to guess how the other players are doing.

For all experiments except for human/naive the following section was included:

Risk neutral bidding strategy

It is common that people lose money in common value actions like this one. Economists hypothesize that this is caused by a combination of two factors: (1) the bidders all have different imprecise estimations of the real common value and (2) the bidders do not take the behavior of the other players into account.

The winner of each auction is the one who submits the highest bid; this bid is based on the signal and the error. At the moment that the winner is revealed, you can be sure that none of the other players was willing to pay the amount the winner bid, or they would have made higher bid. The fact that all other players were not willing to bid as much as the winner did is information that can be taken into account.

Again, it can mean one of two things. (1) It can mean that the signals of all other players are lower than the winner's signal, or (2) it can mean that they play by a different strategy. Of course, it also could be a combination of the two, but it is more likely that some, if not all, of the other players' signals are lower than the winner's signal. Thus, it is very possible that the winner's signal is an overestimation.

Because personal signals are imprecise, bidding at or above your signal risks overpaying, thereby losing money on an auction. According to economic theory, you can protect yourself from the risk of overpaying (and losing money) by bidding no more than your signal minus the general error term (e).

Example:

Your estimate	\$14.75
Error (e)	\$2.00
Risk Neutral bid	\$12.75

It is easy to see why this is the safe strategy: if your signal is the true value plus the maximum error, then subtracting the error will leave you to bid the exact true value. So, in the worst case, having maximum positive error, you will break even. And in all other cases you will make money.

Appendix B: Derivation of probability of winning as a function of error

If all other bidders (j) employ the RNNE bid strategy or any other bidding strategy that is lineary dependent on x_j, the the probability of winning the auction for player i is given by

$$\begin{aligned}
 P_i &= \int p(x_j < x_i | x_0)^{n-1} p(x_0) dx_0 \\
 &= \frac{1}{(2\varepsilon)^n} \int_{x_i-\varepsilon}^{x_i+\varepsilon} x_i + \varepsilon - x_0 dx_0
 \end{aligned}$$

for $x_i \in [x_L + \varepsilon, x_H - \varepsilon]$. If player i unilaterally increases her bid by an amount ρ above RNNE then this increases the probability of winning by different amount depending on the relative values of x_i and x_0 . In particular, if x_i is a high estimate (i.e. $x_i \approx x_0 + \varepsilon$ then increasing one's bid has only a small effect on the probability of winning. However, if x_i is a low estimate (i.e. $x_i \approx x_0 - \varepsilon$ then the probability of winning can be increased from close to zero to 1 if ρ is large enough. To examine the dependence of the probability of winning on ρ , first note that the probability of winning is 1 anytime $\rho \geq x_0 + \varepsilon - x_i$, or $x_0 \leq x_i + \rho - \varepsilon$. This gives the solution to the integral above for any $\rho \in [0, 2\varepsilon]$:

$$\begin{aligned}
 P_i &= \int_{x_i - \varepsilon}^{x_i + \rho - \varepsilon} \frac{1}{2\varepsilon} dx_0 + \\
 &\quad \int_{x_i + \rho - \varepsilon}^{x_i + \varepsilon} \frac{1}{(2\varepsilon)^n} (\rho + x_i + \varepsilon - x_0)^{n-1} dx_0 \\
 &= \frac{\rho}{2\varepsilon} + \frac{1}{n} \left[1 - \left(\frac{\rho}{2\varepsilon} \right)^n \right].
 \end{aligned}$$

For any ρ , note that its effect on the probability of winning an auction diminishes as ε increases. However, if ρ is increased as a function of ε (e.g., $\rho = \alpha\varepsilon$) then the chances of winning will be increased independent of the specifics of the auction.