Can Relaxation of Beliefs Rationalize the Winner's Curse?: An Experimental Study*

Asen Ivanov†
Virginia Commonwealth University

Dan Levin‡
The Ohio State University

Muriel Niederle§
Stanford University

August 27, 2008

Abstract

We use a second-price common-value auction, the *maximal game*, to experimentally study whether the *Winner's Curse* (WC) can be explained by models which retain best-response behavior but allow for inconsistent beliefs. In the maximal game, the WC can be rationalized only by a belief that others use weakly-dominated strategies. Yet, we find that the WC is widespread. In addition, we create environments where, regardless of beliefs, there should be a correction of the WC. We find little evidence of such a correction. Overall, our study suggests that the WC, at least in initial periods of play in the laboratory, represents a more fundamental departure from NE than a mere inconsistency of beliefs.

Keywords: common-value auctions, winner's curse, beliefs, cursed equilibrium, level-k model, quantal response equilibrium

*We would like to thank David Harless and Oleg Korenok for useful discussions.
†301 W. Main Street, room 3149, Richmond, VA 23284; aivanov@vcu.edu
‡1945 N High Street, Arps Hall 410, Columbus, OH 43210; levin.36@osu.edu
§579 Serra Mall, Stanford, CA 94305-6072; niederle@stanford.edu

1 Introduction

In a game of incomplete information, (Bayesian) Nash Equilibrium (NE) has three components: (a) players form beliefs about the type-contingent strategies of the other players, (b) players best-respond to these beliefs, and (c) these beliefs are consistent with actual behavior. Common-value auctions[1] are an important example of games of incomplete information in which behavior, at least in the laboratory, systematically deviates from NE so that at least one of the components (a)-(c) must be violated.

Experimental evidence starting with Bazerman and Samuelson (1983) and Kagel and Levin (1986), as well as many replications by Kagel and Levin and others[2] have documented a *Winner's Curse* (WC) phenomenon in common-value auctions - a systematic overbidding relative to NE which results in massive losses in the lab.[3] Two recent papers, Eyster and Rabin (2005) and Crawford and Iriberri (2007), attempt to explain the WC through theoretical models which retain the components of NE that players are forming and best-responding to beliefs about others' behavior (we shall refer to such models as belief-based) but relax the requirement of consistency of beliefs. Eyster and Rabin introduce the concept of Cursed Equilibrium (CE) in which players fail to fully realize the connection between other players' types and bids and, as a result, succumb to the WC.[4] Crawford and Iriberry use the level-k model of behavior[5] in which level-0 (L_0) players bid in some pre-specified way and level-k (L_k) players ($k = 1, 2, \ldots$) best-respond to a belief that others are L_{k-1}.[6] Ultimately, the success or failure of these theories should be judged by their ability to reconcile deviations from NE for a large body of data from different experimental environments.

Charness and Levin (forthcoming) find that the WC is alive and well in an

[1]In a common-value auction, the value of the object is the same ex post to all bidders but is unknown at the time that bidding takes place. Each bidder receives a private signal which is correlated with the common value.

[2]For example, see Kagel, Harstad, and Levin (1987); Dyer, Kagel, and Levin (1989); Lind and Plott (1991); and the papers surveyed in Kagel (1995, Section II) and Kagel and Levin (2002).

[3]There are also claims of the WC in the field. See Hendricks, Porter, and Boudreau (1987); Hendricks and Porter (1988); and the papers surveyed in McAfee and McMillan (1987, Section XII), Thaler (1988), Wilson (1992, Section 9.2), and Laffont (1997, Section 3).

[4]Carrillo and Palfrey (2008) use CE to explain the occurrence of trade in an experimental design in which a no-trade theorem holds.

[5]This model was first suggested by Stahl and Wilson (1995) and Nagel (1995).

[6]CE and the level-k model can be applied to environments other than common-value auctions.

individual-choice variant of the "acquiring a company" game, i.e. in a situation where the WC cannot be explained by incorrect beliefs about others' behavior.[7] In the current paper, we investigate experimentally whether belief-based models, such as CE and the level-k model, can explain the WC in an actual auction.[8] We focus on initial periods of play as this seems like a natural starting point for evaluating belief-based theories.

In our experiment, we use a particular second-price common-value auction, called the *maximal game*, where actual behavior is likely to exhibit the WC. In the maximal game, each player receives an *i.i.d.* signal and the common value of the object being auctioned equals the highest signal. The object goes to the highest bidder who pays the second highest bid. With signals in $[0, 10]$, a player with signal x, can make four qualitatively different kinds of bids, b:

(i) $b < x$ - we call this underbidding;

(ii) $b = x$ - we call this bidding one's signal;

(iii) $x < b \leq 10$ - we call this overbidding;

(iv) $b > 10$ - we call this bidding above 10.

$b < x$ is weakly dominated (by $b = x$) since, under the second-price rule, one may lose the auction at a price below one's signal even though the value of the object is greater than or equal to x. $b > 10$ is also weakly dominated but, unlike $b < 10$, may result in a negative payoff. Because they are weakly dominated, $b < x$ and $b > 10$ can hardly be explained by belief-based theories.

If a player believes that others will never underbid (because it's weakly dominated), $x < b \leq 10$ turns out to be weakly dominated as well. Hence, the game is two-step dominance solvable and bidding one's signal is the unique remaining strategy.[9]

[7]Although not a common-value auction, the "acquiring a company" game shares the informational features of common-value auctions so that a WC phenomenon can arise.

[8]Pevnitskaya (2008) investigates whether inconsistent beliefs can explain behavior in private value auctions.

[9]Note that for bidding one's signal to be a best response, it suffices that a player believes that others never make a weakly dominated bid, without any need for common knowledge of rationality.

The most interesting case is $x < b \leq 10$. This case constitutes a WC - subjects are bidding above NE and, to the extent that other subjects are doing the same, may lose money.

The maximal game is particularly suitable for testing whether the WC can be explained by belief-based models. In particular, because the game is two-step dominance-solvable, $x < b \leq 10$ is weakly dominated (by $b = x$) unless a player believes that others are playing weakly-dominated strategies, i.e. are underbidding (at least for some signals). This puts a strain on belief-based models as possible explanations for overbidding. Therefore, to the extent that we observe the WC in the maximal game, belief-based models are less plausible explanations than in other common-value auctions.

It is nevertheless possible that subjects overbid because they believe that others are underbidding. For example, we will show that this is the case for players in a CE as well as for L_1 players in the level-k model. Our treatments provide a multipronged approach to testing this possibility. This approach does not rely on belief elicitation which raises issues of truthful reporting and may change subjects' cognitive processes.

In part I of our *Baseline* treatment, we let each subject bid in 11 two-bidder auctions (without any feedback). In part II, each subject participates in another 11 auctions. However, this time she bids against the computer which uses *her bid function from part I.* If subjects (for all signals) overbid in part I, then their best-response in part II will be to start bidding $b = x$, or at least to shift their bids downwards. If we see a downward shift of bids in part II, it suggests that subjects are being strategic by responding to the behavior of the other player, at least when this behavior is made salient through the fact that it is their own behavior from part I. Such a change in behavior in part II would leave hope for belief-based models. On the other hand, if overbidding is not corrected downwards in part II, then this seriously undermines any belief-based explanations of overbidding in part I.

In part I of the *Baseline* treatment, we find that the largest proportion of bids is of the form $x < b \leq 10$ so that the WC is alive and well. More importantly, of those subjects who (predominantly) overbid in part I, only a minority switch in part II to (predominantly) bidding $b = x$ or $b < x$. The majority continue to (predominantly) overbid in part II without any tendency for a downward correction of bids. These results continue to hold when we explicitly show subjects their bid functions from part I in part II (which we do in our *ShowBidFn* treatment).

In our *MinBid* treatment, we take a different approach to testing the validity of belief-based explanations of the WC. This treatment differs from the *Baseline* treatment only in that subjects are explicitly not allowed to underbid. In the *MinBid* treatment, overbidding is weakly dominated and can hardly be explained by belief-based models. Despite that, the frequency of overbidding in part I is similar to the frequency of overbidding in the *Baseline* and *ShowBidFn* treatments. Moreover, the (average) magnitude of overbidding in part I of the *MinBid* treatment, on the one hand, and of the *Baseline* and *ShowBidFn* treatments, on the other, is practically identical.

In addition, in all three treatments, $b > 10$ is not uncommon. Such bids fall prey to the WC, but because they are weakly dominated, they can hardly be explained by belief-based models.

Our experimental results cast serious doubts on any belief-based explanations of the WC in initial periods of play in the maximal game. These results, together with Charness and Levin (forthcoming), strongly suggest that the WC, at least in initial periods of play in the laboratory, represents a more fundamental departure from NE than a mere inconsistency of beliefs. More generally, our results cast doubt on subjects forming beliefs and best-responding to them in initial periods of play in games of incomplete information.

Finally, we also consider whether Quantal Response Equilibrium (McKelvey and Palfrey (1995)) could explain behavior in our experiment. QRE relaxes the requirement for best-responding rather than the consistency of beliefs, and the maximal game is not specifically designed to test it. Nevertheless, we find evidence which is at odds with QRE.

We proceed as follows. In section 2, we present the maximal game and derive the relevant theoretical predictions. In section 3, we describe our experimental design and in section 4 we examine the experimental data. In section 5 we consider QRE. Section 6 concludes.

2 Theoretical Considerations

We begin by describing the maximal game. There are n bidders, each of which privately observes a signal X_i that is *i.i.d.* from a distribution $F(\cdot)$ on $[0, 10]$. We make no assumptions on $F(\cdot)$. Let $X^{\max} = max(\{X_i\}_{i=1}^{n})$ be the highest of the n signals. Let x_i and x^{max} denote particular realizations of X_i and X^{max}, respectively. Given $(x_1, \ldots, x_n)$, the *ex-post* common valuation of the bidders in the maximal game is $v(x_1, \ldots, x_n) = x^{\max}$.

Bidders bid in a second-price auction where the highest bidder wins, earns the common-value, $x^{\max}$, and pays the second highest bid. In case of a tie, the winning bidder is determined randomly with all tying bidders getting the object with equal probability.

Let $b_i(\cdot)$ denote player i's bid function. Also, let x_{-i} and b_{-i} be particular profiles of signals and bids, respectively, for players other than i. Given signal x_i, a bid b is weakly dominated for player i *iff* there is a bid b', such that (i) for any x_{-i} and b_{-i}, bidding b' is no worse than bidding b, and (ii) for some x_{-i} and b_{-i}, bidding b' is strictly better than bidding b. A bid function, $b_i(\cdot)$, is weakly dominated *iff* for some x_i, $b_i(x_i)$ is a weakly dominated bid. We can now state our first result.

Proposition 1 *$b(x_i) = x_i$ is the unique bid function remaining after two rounds of iterated deletion of weakly dominated bid functions. In the first round, all bid functions $b_i(\cdot)$, such that $b_i(x_i) < x_i$ for some x_i, are deleted. In the second round, all bid functions $b_i(\cdot)$, such that $b_i(x_i) > x_i$ for some x_i, are deleted.*

The proof is in the appendix. Here, we give the intuition. Underbidding is weakly dominated since, under the second-price rule, one could lose the auction at a price below one's signal even though the value of the object is greater than or equal to one's signal. Given that no one underbids, $b_i(x_i) > x_i$ is weakly dominated for any x_i, because, in case the highest bid among others is between x_i and $b_i(x_i)$, i makes non-positive (and possibly negative) profits.

That bidding one's signal is a NE, follows directly from proposition 1. In fact, we can say more than that (the proof is in the appendix)[10]:

[10] Under standard assumptions on $F(\cdot)$, we could simply invoke proposition 1 in Pesendorfer and Swinkels (1997), so that no proof would be necessary. However, these assumptions do not hold in the case of the discrete distribution in our experiment.

Proposition 2 *The bid function $b(x_i) = x_i$ is the unique symmetric NE (including mixed strategies).*

We have shown that overbidding is weakly dominated if one believes that others do not use weakly dominated bid functions and that overbidding cannot be part of a symmetric NE.[11] However, overbidding can arise within the level-k model and within CE.

First, let us consider the level-k model. In this model, level-0 (L_0) players bid in some pre-specified way and level-k (L_k) players ($k = 1, 2, \ldots$) best-respond to a belief that others are L_{k-1}. For normal-form games with finite actions in the literature, L_0 is modeled as choosing each action with equal probability. For auction settings, Crawford and Iriberry (2007) have two versions of L_0. The *Random* L_0 (RL_0), regardless of its signal, bids uniformly over all bids between the minimal and maximal value of the object (i.e. in $[0, 10]$ in our settings), and the *Truthful* L_0 (TL_0) bids its signal. RL_k/TL_k ($k \geq 1$) best-respond to RL_{k-1}/TL_{k-1}. Below, we show that TL_1 and RL_1 can bid above their signals. First, let us consider TL_1.

Proposition 3 *TL_1 can use any bid function $b^{TL_1}(\cdot)$, with $b^{TL_1}(x_i) \geq x_i$, $\forall x_i$.*

The proof is in the appendix.[12] It rests on the fact that bidding above one's signal cannot lead to negative profits, since the second highest price cannot be higher than x^{max} (others are bidding their signals).[13] Now, let us consider RL_1.

Proposition 4 *The bid function of RL_1 is $b^{RL_1}(x_i) = E(X^{max}|X_i = x_i) \geq x_i$. If $F(x_i) < 1$, the inequality is strict.*[14]

The proof is in the appendix. It hinges on the fact that, because an RL_0's bid is uninformative about its signal, RL_1 cannot draw any inference about X^{max} from

[11] In our experiment, matching of subjects is anonymous and there is no feedback, so that asymmetric Nash Equilibria do not seem plausible.

[12] This proposition implies than, when all others are playing the symmetric NE, player i can bid anything greater than or equal to her signal.

[13] Because the behavior of a TL_1 is not uniquely determined, neither is the behavior of a TL_k for $k \geq 2$. The point, however, is that a TL_1 can rationalize overbidding (and even bidding above 10).

[14] If signals have the discrete uniform distribution on the set $\{0, 1, 2, \ldots, 10\}$ and there are two bidders (this is relevant for our experiment), then $b^{RL_1}(x_i) = E(X^{max}|X_i = x_i) = \frac{x_i^2 + x_i + 110}{22}$. This is greater than x_i $\forall x_i \in [0, 10]$, strictly so if $x_i < 10$.

winning the auction.[15]

Let us turn to Eyster and Rabin's CE. In a *chi*-CE ($\chi \in [0,1]$), players are assumed to best respond to other players' behavior in a certain sense. Each player i believes that with probability χ each other player j chooses a bid that is type-independent, and which is distributed according to the ex ante distribution of player j's bids. In addition, players believe that with probability $1-\chi$ each other player j chooses a bid according to player j's actual type-dependent bid function. Thus, χ captures players' level of "cursedness": if $\chi = 0$, we have a standard Bayesian Nash equilibrium, and if $\chi = 1$, players draw no inferences about other players' types.[16] Based on proposition 5 in Eyster and Rabin (2005), we can state:

Proposition 5 *Assuming that X_i has a strictly positive pdf, the following bid function constitutes a symmetric χ-CE:*

$$b^{CE}(x_i) = (1-\chi)x_i + \chi E(X^{max}|X_i = x_i)$$

Note that when $\chi = 0$, the proposed CE reduces to Milgrom and Weber's (1982) equilibrium for second-price common-value auctions. When $\chi = 1$, players bid $E(X^{max}|X_i = x_i)$ because, just like RL_1 players, they draw no inference from winning the auction about others' signals. In general, the second term in $b^{CE}(\cdot)$ represents the fact that players in a χ-CE underappreciate the information content of winning.

Note that $b^{CE}(x_i) \geq x_i$ with strict inequality whenever $\chi > 0$ and $F(x_i) < 1$. Thus, CE can also rationalize overbidding.[17]

[15] A RL_2 can use any bid function $b(\cdot)$, such that $b(x_i) < E(X^{max}|X_j = 0)$ for all $x_i < E(X^{max}|X_j = 0)$ and $b(x_i) = x_i$ for all $x_i \geq E(X^{max}|X_j = 0)$. Because the behavior of RL_2 is not uniquely determined, neither is the behavior of a RL_k for $k \geq 3$. The point, however, is that a RL_1 can rationalize bidding above the signal.

[16] There are versions of CE in which each player has a different χ (and thus a different level of cursedness). Although, we could derive predictions for the maximal game in such cases, our point is merely to show that CE, even without allowing for subject-specific χ, can lead to overbidding.

[17] Although the assumption of a strictly positive pdf is not satisfied for the discrete distribution in our experiment, we suspect that the proposition nevertheless holds. However, we did not go as far as formally proving this because our data does not support belief-based explanations of overbidding. Therefore, we run little danger of incorrectly explaining overbidding through CE.

3 Experimental Design

3.1 Treatments

The experiment consists of three treatments: the *Baseline*, *ShowBidFn*, and *MinBid* treatments, which we now describe.

The *Baseline* treatment consists of two parts. In part I, subjects play the maximal game for 11 periods. In each period, subjects are randomly and anonymously rematched in separate two-player auctions. Each subject's signals for the 11 auctions are drawn without replacement from the set $\{0, 1, 2, \ldots, 10\}$.[18] In a given auction, each signal which was not received by the subject in a previous auction has an equal chance of being drawn and is independent of the other bidder's signal. Subjects can bid anything between 0 ECU and 1000000 ECU.[19] Subjects receive no feedback whatsoever until the experiment is over. This minimizes any effects from learning. It also guarantees that, in any auction, each bidder's prior over the other bidder's signal is the discrete uniform distribution on the set $\{0, 1, 2, \ldots, 10\}$ (since no subject sees any other subject's past signals).

Part II is similar to part I. The only difference is that each subject i bids against the computer rather than against another subject. The computer, which "receives" a signal just like any human bidder, mimics i's behavior from part I by using the same bid function that i used in part I. In particular, if the computer receives signal y, it makes the same bid that i made in part I when she received signal y. In effect, in part II each subject is playing against herself from part I.

Consider a subject i who overbids (for all signals) in part I.[20] From proposition 1, it follows that bidding her signal is a best-response in part II. Underbidding may not be a best-response, but it is at least a response in the right direction. If we see i start bidding her signal or underbidding in part II, this would suggest she is (best-)responding to a belief about the behavior of the other player, at least when the

[18]Our design for part I ensures that each subject will receive each of the eleven signals from the set $\{0, 1, 2, \ldots, 10\}$ exactly once. In effect, we are eliciting subjects' bid functions. This simplifies the design of part II.

[19]We thought about not allowing bids above 10. However, this seemed artificial so we opted for no such restriction.

[20]To be precise, overbidding is not possible for signal 10: a subject can either underbid, bid her signal, or bid above 10. Therefore, the correct statement is "a subject i who overbids for all signals 0-9 and bids above 10 for signal 10".

other player is herself from part I so that the correct belief is obvious. This would at least be consistent with belief-based models.

What if in part II subject i continues overbidding, or even starts bidding above 10? Note that, even though bidding her signal is a best-response, continuing to overbid in part II could also be a best-response.[21] However, for this to be the case, i needs to shift her bid function in part II, $b_i^{II}(\cdot)$, downwards in a way that, for all signals x_i, none of the bids she made in part I lies in $(x_i, b_i^{II}(x_i)]$.[22] Therefore, if i starts bidding above 10 in part II, we can be quite sure she is not best-responding to her behavior from part I. If i continues overbidding in part II, we would be sceptical that she is best-responding to her behavior from part I - after all, why not simply bid her signal instead of shifting her bid function down in a seemingly complicated way. Nevertheless, in our data analysis we will need to verify that any continued overbidding in part II is not part of a best-response.

Say, subject i continues overbidding in part II without best-responding. It is still important to check whether she corrects her overbidding from part I downwards in part II. Even if the size of such a correction is not optimal, any downward shift in bids would suggest that i is at least responding in the right direction to her behavior from part I. Therefore, in our analysis, we will look not only at whether overbidding persists in part II, but also at whether there is a downward correction.

What can we conclude about belief-based models if overbidding persists without any downward correction in part II so that subjects don't seem to be (best-)responding to their behavior from part I? Belief-based models, such as the level-k model and CE, could explain overbidding in part I by relaxing the requirement that beliefs are correct. However, such models are much less plausible explanations for overbidding in part II given that subjects in part II are unlikely to have incorrect beliefs about their own bidding behavior in part I (which was just a few minutes ago).[23] Therefore, if in part II overbidding persists without any downward correction even though belief-based models are unlikely explanations, then there is little reason to think that overbidding in part I is driven by beliefs.

[21]For example, if i bids 10 for all signals in part I, bidding $5 + \frac{x_i}{2}$ is a best response in part II.

[22]Otherwise, there's a positive probability she wins the auction and loses money.

[23]For example, to invoke the level-k model, one has to assume that subjects erroneously believe in part II that they played like a TL_0/RL_0 in part I. To invoke cursed beliefs, one has to assume that subjects fail to fully recognize in part II that they bid differently for different signals in part I.

One could still argue that for some reason subjects in part II simply ignore their behavior from part I or perhaps are unable to recall it[24], and instead form TL_1/RL_1 or cursed beliefs. In order to stress-test whether overbidding persists without any downward correction in part II, we introduce our *ShowBidFn* treatment. This treatment is the same as the *Baseline* treatment with the sole difference that in part II we explicitly show subjects their bid functions from part I. In this treatment, it seems even more implausible that a subject in part II should have incorrect beliefs about her own behavior in part I.

Our *MinBid* treatment provides yet another test of the persistence of overbidding in a context where belief-based explanations seem implausible. This treatment is the same as the *Baseline* treatment with the sole difference that subjects are not allowed to underbid. Even in part I of this treatment, any bid above one's signal is weakly dominated by bidding one's signal. Hence overbidding can hardly be explained by belief-based models.[25]

The *MinBid* treatment is very useful because we do not need to rely on indirect evidence from part II to conclude that overbidding in part I is not belief-driven. However, it is possible that, because bidding one's signal is on the boundary of the set of admissible bids, this may bias subjects towards bidding above their signals. Note, however, that if this bias is driving overbidding in part I of the *MinBid* treatment and belief-based models are driving overbidding in part I of the *Baseline* and the *ShowBidFn* treatment, then there is no reason that the frequency and (average) magnitude of overbidding should be similar in the *MinBid* treatment and in the other two treatments.

3.2 Procedures

We conducted three sessions of the *Baseline* (62 subjects), two sessions of the *ShowBidFn* (46 subjects) and one session of the *MinBid* treatment (26 subjects).[26]

Subjects in the experiment were students at The Ohio State University who were

[24]Note that a subject who overbid (for all signals) in part I need not remember her exact bid function from part I. It suffices if she remembers that she overbid (for all signals) in part I.

[25]A TL_1 could still overbid, but unlike in part I of the *Baseline* or the *ShowBidFn* treatment, this would be weakly dominated.

[26]We also conducted a pilot session for the *Baseline* treatment (26 subjects).

enrolled in undergraduate Economics classes. The sessions were held at the Experimental Economics Lab at OSU and lasted around 45 minutes. At the start of each session, the experimenter read the instructions for part I aloud as subjects read along, seated at their computer terminals. After that, subjects did a practice quiz. Experimenters walked around the room checking subjects' quizzes, answering questions and explaining mistakes.[27] After we conducted part I of the relevant treatment, we went over the instructions for part II. After part II, subjects were paid. Subjects' earnings consisted of a \$5 show-up fee, plus 10 ECU starting balances, plus their cumulative earnings from the 22 auctions[28], converted at a rate of \$0.5 per ECU. Average earnings were \$18.53, \$18.03, and \$15.53 in the *Baseline*, *ShowBidFn*, and *MinBid* treatment, respectively. The instructions for the *Baseline* treatment are in the appendix.[29] In the appendix, we also provide a printout from the screen of a subject in part II of the *ShowBidFn* treatment. The screen in part I of all three treatments and part II of the *Baseline* and *MinBid* treatments is exactly the same, but without the table showing the subject's bid function from part I. The experiment was programmed and conducted with the software z-Tree (Fischbacher (2007)).

4 Results

We start by studying behavior in parts I and II within each treatment. We first place each bid b, given a subject's signal x, in one of the following four categories: (i) $b < x - 0.25$, (ii) $x - 0.25 \leq b \leq x + 0.25$ which we simply denote by $b \sim x$, (iii) $x + 0.25 < b \leq 10$, and (iv) $b > 10$[30]. That is, we count all bids within 0.25 ECU of one's signal as if they were precisely equal to the signal, thus avoiding too strict an interpretation of bidding one's signal.[31] Based on this, we look at the percentages of bids that fall in each category and at how these percentages change from part I to

[27]We had a strong sense that the practice quiz was very useful for ensuring that subjects understood the task.

[28]In case a subject made losses which could not be covered by the 10 ECU starting balances, she was paid just her \$5 show-up fee.

[29]The instructions in the other two treatments are very similar and are available upon request from the authors.

[30]Actually, for signal x=10, a bid needs to be above 10.25 in order to fall in category (iv); a bid $9.75 \leq b \leq 10.25$ falls in category (ii). We ignore this in our notation.

[31]Counting only bids which are precisely equal to the signal in category (ii) (and adjusting the other categories appropriately) does not change any of our results.

	$b < x - 0.25$	$b \sim x$	$x + 0.25 < b \leq 10$	$b > 10$
Part I	16.1%	21.1%	41.6%	21.1%
Part II	14.2%	26.1%	37.2%	22.4%

Table 1: Percentage of bids in each category for the *Baseline* treatment.

part II. Extending our analysis at the individual level, we classify subjects according to the category in which their bids predominantly fall, and we study how subjects' behavior changes from part I to part II. After that, we turn to the question of how behavior in part I compares between the *Baseline* and *ShowBidFn* treatments, on the one hand, and the *MinBid* treatment, on the other.

4.1 Percentage of Bids in each Category

Table 1 shows the percentage of bids in each category in part I and part II in the *Baseline* treatment. Based on the table, we can state our first result.

Result 1

(1) In part I, the largest percentage of bids are of the form $x+0.25 < b \leq 10$ *(41.6%).*

(2) In part I, a considerable percentage of bids are weakly dominated, i.e. they are of the form $b < x - 0.25$ *or* $b > 10$ *(37.2%).*

(3) In part II, the largest percentage of bids remain of the form $x < b \leq 10$ *(37.2%).*

Point (1) shows that overbidding is widespread. Point (2) shows that subjects are even making weakly dominated bids quite often. Point (3) suggests that overbidding largely persists in part II.

Tables 5 and 6 in the appendix are the analogues of table 1 for the *ShowBidFn* and *MinBid* treatment, respectively. As can be seen from the tables, result 1 also holds for these two treatments.[32]

[32] In the *MinBid* treatment, weakly dominated bids occur less frequently, largely because subjects cannot underbid.

Part I / II	*Underbidders*	*Signal-bidders*	*Overbidders*	*Above-10-bidders*	*Indeterminate*	
Underbidders	2	0	2	1	0	**5**
Signal-bidders	0	5	3	1	0	**9**
Overbidders	1	5	14	1	4	**25**
Above-10-bidders	2	1	1	6	0	**10**
Indeterminate	2	2	3	5	1	**13**
	7	**13**	**23**	**14**	**5**	

Table 2: Transition table for the *Baseline* treatment.

4.2 Subject Classification

Table 1 does not reveal whether all subjects make bids in each category with the same frequency or whether each subject bids predominantly in one of the four categories. It also does not reveal whether the subjects who overbid in part I are the same ones who overbid (or even bid above 10) in part II.[33] To shed light on these issues, we classify subjects (separately for each part of each treatment) in the following way: *Underbidders*/*Signal-bidders*/*Overbidders*/*Above-10-bidders* are those who make bids of the form $b < x - 0.25/b \sim x/x + 0.25 < b \leq 10/b > 10$ in at least 6 (out of 11) auctions; subjects who are neither of the preceding are classified as *Indeterminate*.[34]

In light of our discussion in section 3, note that if *Overbidders* from part I become *Signal-bidders* or *Underbidders* in part II, this would suggest that they are (best-)responding to their behavior from part I. If, on the other hand, they remain *Overbidders* or even become *Above-10-bidders* in part II, this would cast doubt on any best-response behavior in part II.

Table 2 shows, for the *Baseline* treatment, how many subjects were in each class in part I (last column) and part II (last row). The table also shows how subjects switched between classes from part I to part II. For example, the entry in the first row and the third column shows that 2 subjects who were *Underbidders* in part I became *Overbidders* in part II. Based on the table, we can state:

[33] It could be that overbidding in part II comes from subjects who are best-responding to the fact that they underbid in part I.

[34] We also conducted our analysis by using 7 or 8 (instead of 6) class-consistent decisions as the cutoff for a player to be assigned to a class. This didn't affect the analysis much apart from increasing the number of *Indeterminate* subjects.

Result 2

(1) In part I, Overbidders are the largest class (40.3%).

(2) In part I, a large percentage of subjects make a weakly dominated bid ($b < x - 0.25$ or $b > 10$) in at least 6 (out of 11) auctions (30.7%).[35]

(3) Only a minority of Overbidders from part I become Signal-bidders or Underbidders in part II (24%).

(4) The majority of Overbidders from part I either remain Overbidders or even become Above-10-bidders in part II (60%).

Result 2 is in line with result 1. In addition, it shows that only a minority of *Overbidders* from part I (best-)respond in part II by becoming *Signal-bidders* or *Underbidders*.[36] Instead, the majority of them remain *Overbidders* or even become *Above-10-bidders*.

Tables 7 and 8 in the appendix are the analogues of table 2 for the *ShowBidFn* and *MinBid* treatment, respectively. As can be seen from the tables, result 2 also holds for these two treatments.[37]

It is fairly clear that the *Overbidder* from part I who becomes an *Above-10-bidder* in part II is not best-responding in part II.[38] However, we need to make sure that subjects who are *Overbidders* in parts I and II are indeed not best-responding in part II. We would also like to check whether they are at least responding in the right direction by lowering their bids downwards.

The second column in table 3 shows, for each subject who was an *Overbidder* in parts I and II of the *Baseline* treatment, what percent of bids in part II is a

[35] This percentage includes all *Underbidders* and *Overbidders*, as well as 4 *Indeterminate* subjects.

[36] The one *Overbidder* from part I who becomes an *Underbidder* in part II foregoes 7.46 ECU in expected profits by not behaving optimally in part II. The five *Overbidders* from part I who become *Signal-bidders* in part II forego on average only 0.53 ECU in expected profits in part II.

[37] In the *ShowBidFn*/*MinBid* treatment the percentage of subjects who make a weakly dominated bid in at least 6 (out of 11) auctions is 28.3%/7.8%. In the *MinBid* treatment, the percentage is smaller largely because subjects cannot underbid.

[38] In fact, she foregoes 29.64 ECU in expected profits by not behaving optimally in part II.

Subject ID	% of bids in part II which are a best-response	foregone expected profits (in ECU) in part II	% of $x + 0.25 < b \leq 10$ from part I reduced in part II	% of $x + 0.25 < b \leq 10$ from part I increased in part II
28	45%	0.27	40%	0%
29	9%	2.77	20%	50%
38	36%	1.23	43%	57%
46	27%	6.23	14%	71%
55	18%	3.00	75%	13%
57	36%	2.59	29%	71%
59	9%	15.95	13%	75%
66	9%	4.05	20%	60%
75	9%	5.00	67%	33%
76	9%	6.00	38%	50%
80	18%	7.55	29%	57%
82	45%	2.68	57%	14%
84	9%	17.32	0%	88%
86	36%	4.09	50%	0%
Mean	**23%**	**5.62**	**35%**	**46%**

Table 3: Best-response behavior in part II for subjects who are *Overbidders* in parts I and II in the *Baseline* treatment.

best-response to part I behavior. The third column shows how much each subject is foregoing in expected profits from not bidding optimally in part II. From the table, we see that subjects are seldom best-responding and are foregoing, on average, 5.62 ECU in expected profits in part II. This is 15% of average earnings in the *Baseline* treatment. The fourth/fifth column of table 3 shows for what percent of bids of the form $x+0.25 < b \leq 10$ from part I, the corresponding bid in part II (i.e. the one made for the same signal) is strictly lower/higher. As can be seen from the table, subjects are, on average, even more likely to increase a bid of the form $x + 0.25 < b \leq 10$ from part I in part II than they are to reduce it.

Figure 1 gives us another view of whether subjects who are *Overbidders* in parts I and II correct their bidding downwards in part II. Here, we plot, for each signal, the median bid in part I (circles) and part II (stars).[39] Based on the figure, we see no downward correction of bids in part II. Let us summarize our findings regarding the (best-)response behavior in part II of subjects who are *Overbidders* in parts I and II:

[39]We plot median, rather than average, bids because averages are distorted by bids above 10 (which are sometimes very high, even up to 1000000).

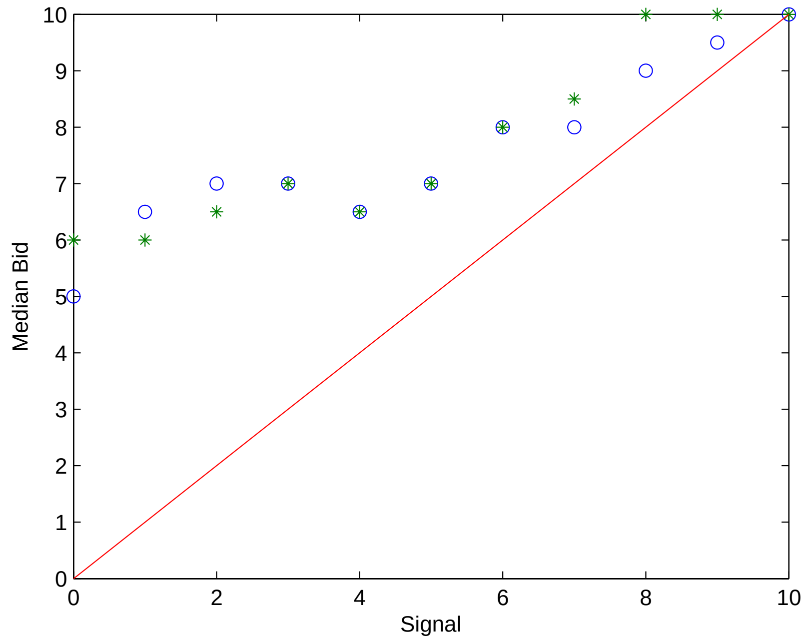

Figure 1: Median bids in parts I (circles) and II (stars) for subjects who are *Overbidders* in parts I and II of the *Baseline* treatment.

Result 3 *For subjects who are Overbidders in parts I and II, we find that:*

(1) In part II, they forego, on average, substantial expected profits (5.62 ECU).

(2) In part II, there is no evidence of a downward correction of bids.

Tables 9 and 10 in the appendix are analogous to table 3 for the *ShowBidFn* and *MinBid* treatment, respectively; figures 4 and 5 are analogous to figure 3 for the *ShowBidFn* and *MinBid* treatment, respectively. Based on these tables and figures, we see that result 3 also holds for the *ShowBidFn* and *MinBid* treatments.

4.3 *Baseline* and *ShowBidFn* vs. *MinBid*

If overbidding in part I of the *Baseline* and *ShowBidFn* treatments is driven by a belief that others are underbidding, we would expect a decrease in the frequency and/or (average) magnitude of overbidding in part I of the *MinBid* treatment.

	$b < x - 0.25$	$b \sim x$	$x + 0.25 < b \leq 10$	$b > 10$
Baseline and *ShowBidFn*	16.5%	22%	42.8%	18.7%
MinBid	0%	28.3%	60.5%	11.2%

Table 4: Percentage of bids in each category for *Baseline* and *ShowBidFn* vs. *MinBid*.

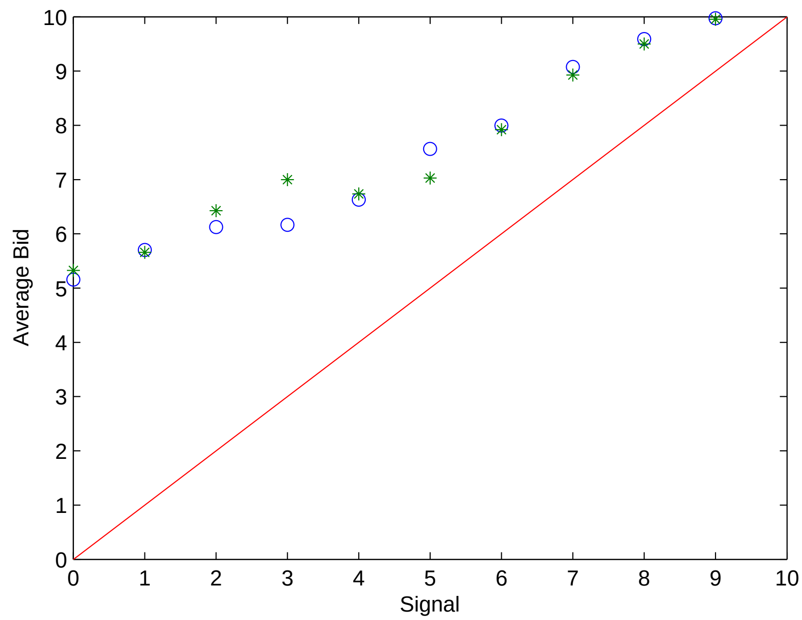

Figure 2: Average bids in part I of *Baseline* and *ShowBidFn* (circles) and *MinBid* (stars) (based on bids of form $x + 0.25 < b \leq 10$).

Table 4 shows the percentage of bids in each category for the *Baseline* and *ShowBidFn* treatments, on the one hand, and the *MinBid* treatment, on the other.[40] We see that overbidding is in fact even more frequent in the *MinBid* than in the other two treatments. This is probably partially due to the fact that in the *MinBid* treatment underbidding is impossible so that all bids are distributed in three, rather than four, categories. Given this, the frequencies of overbidding seem quite comparable.

[40]We pool the data from part I of the *Baseline* and *ShowBidFn* treatments because part I is the same in both treatments.

How about the magnitude of overbidding? Figure 2 shows, for each signal, the average bid of the form $x + 0.25 < b \leq 10$ in part I of the *Baseline* and *ShowBidFn* treatments (circles) and in part I of the *MinBid* treatment (stars). Average bids are astonishingly close.[41] We can summarize:

Result 4 *Relative to the Baseline and ShowBidFn treatments, we find no evidence in the MinBid treatment of:*

(1) a lower frequency of bids of the form $x + 0.25 < b \leq 10$*;*

(2) a reduction in the average size of bids of the form $x + 0.25 < b \leq 10$*.*

This provides additional evidence that overbidding is not driven by beliefs.

5 Quantal Response Equilibrium

In Quantal Response Equilibrium (McKelvey and Palfrey (1995)), players have correct beliefs about others' behavior, but choose noisy best-responses. The likelihood of a particular error depends on how costly that error is as well as on a precision parameter λ (within the usual logit specification). When $\lambda = 0$, subjects choose randomly; as λ goes to ∞, the probability of choosing a best-response goes to 1.

Our experiment was designed specifically to test belief-based theories in which, unlike in QRE, subjects best-respond without noise (to possibly inconsistent beliefs). Although we were not able to solve for QRE in the maximal game, we can still say something about whether QRE is a plausible explanation of behavior. In particular, for each subject given her own behavior from part I, we test the hypothesis that in part II she is choosing randomly (i.e. has precision $\lambda = 0$) rather than choosing noisy best-responses with $\lambda > 0$. If we fail to reject this hypothesis for many subjects, this suggests that assuming noisy best-responses does little to organize behavior in part II. This would make QRE with $\lambda > 0$ (the only interesting case[42]) a quite unlikely

[41] Not surprisingly, using random effects regressions, we find no statistically significant differences.

[42] It is theoretically elegant that QRE nests random behavior as a special case; however, QRE is an interesting concept only if $\lambda > 0$.

explanation of behavior in part I given that it assumes not only strictly positive precision in best-responding, but also consistency of beliefs.

We can reject the said hypothesis (at the 5% level) for only 39%/26%/46% of subjects in the *Baseline*/*ShowBidFn*/*MinBid* treatment.[43] This is at odds with QRE.

6 Concluding Remarks

A recent paper by Charness and Levin (forthcoming) already challenged the robustness of belief-based explanations of the WC by finding massive WC behavior in an individual-choice variant of the "acquiring a company" game. In the current paper, we investigate whether the WC in initial periods of play can be explained by belief-based models within the context of an actual auction.

Our experiment is based on the maximal game which has the desirable property that it is two-step dominance solvable. This allows us to take a multipronged approach to investigating whether the WC could be explained by belief-based models.

First, overbidding in the maximal game can be rationalized by beliefs only if players believe that others are playing weakly dominated strategies. This already puts a strain on any belief-based explanations of the WC. Despite that, we find that overbidding is widespread.

Second, even if in part I of our treatments subjects are overbidding because they believe that others are playing weakly dominated strategies, they should still start bidding their signals, or at least correct their bids downwards, in part II where they are basically playing against themselves from part I. Only a minority of those subjects who (predominantly) overbid in part I switch in part II to (predominantly) bidding their signals or underbidding. The majority continue to (predominantly) bid above their signals in part II without any downward correction of bids.

These results continue to hold in our *ShowBidFn* treatment in which we explicitly show subjects their bid functions from part I in part II.

Finally, explicitly not allowing subjects to underbid, as we do in our *MinBid*

[43] The hypothesis is tested via a score test. In finding the likelihood of a subject's bids in part II (under the null of $\lambda = 0$), we needed to specify the range of all possible bids. Taking this range to be $[0, 1000000]$ was impractical. Instead, we took this range to be $[0, A]$ and truncated bids in part II at A, where A was set at 10. Setting A=15 slightly decreases the number of subjects for whom we can reject the null.

treatment, makes overbidding weakly dominated. Despite that, overbidding in part I persists in the *MinBid* treatment with a similar frequency and the same (average) magnitude as in the other two treatments.

Overall, our study suggests that the WC, at least in initial periods of play in the laboratory, represents a more fundamental departure from NE than a mere inconsistency of beliefs.

This begs the question: what does explain the WC? In the rich literature on overbidding in private-value and common-value auctions, a line of research uses models with boundedly rational agents (with some limits on cognitive ability) and/or with non-standard preferences. To review this research takes us beyond the scope of this study.[44] Instead, let us conclude with a few speculative remarks.

Kagel, Harstad and Levin (1987) find systematic bidding above values in private-value sealed-bid second-price auctions, even though bidding one's value is a dominant strategy.[45] At the same time, bidders converge almost immediately to the dominant strategy in the strategically equivalent English auction. A plausible explanation of this, given by the authors, is that realizing the dominant strategy in the sealed-bid auction requires non-trivial reasoning through various contingent scenarios (regarding the possible order of one's value, bid, and others' bids). In contrast, the English auction eliminates the need for contingent reasoning: as the clock-price ascends, simply answering correctly the question "should I stay or drop out?" leads to the dominant strategy. In the context of their individual-choice variant of "acquiring a company" game, Charness and Levin (forthcoming) suggest that the "origin" of the WC also lies in an inability of subjects to engage in contingent reasoning.[46] Based on these studies, we conjecture that the WC is primarily driven by an inability to realize that the expected value of the object should be computed contingent on winning. This inability constitutes a departure from NE at a more fundamental level than is present in belief-based models. The latter presuppose that subjects can form a clear enough picture of all possible contingencies in order to form (and best-respond to) a belief. Of course, more research is needed to verify this conjecture.

[44] See in Kagel (1995); and Kagel and Levin (2002, forthcoming).

[45] This result has been replicated several times (see Kagel (1995), Kagel and Levin (forthcoming)).

[46] Ivanov, Levin and Peck (forthcoming) suggest that subjects' capacity to think about various scenarios in the future (subjects' foresight) plays a key role in determining whether subjects wait to learn from others in an endogenous timing investment game.

References

[1] Bazerman, M H and W. F. Samuelson, 1983: "I won the Auction But Don't Want the Prize," *Journal of Conflict Resolution*, 27, 618-34.

[2] Carrillo, Juan D. and Thomas R. Palfrey (2007): "No Trade," University of Southern California working paper, available at http://www-rcf.usc.edu/ juandc/papers-others.html

[3] Charness, G. and D. Levin (forthcoming): "The Origin of the Winner's Curse: A Laboratory Study," *American Economics Journal Microeconomics.*

[4] Crawford, V. P. and N. Iriberri (2007): "Level-k Auctions: Can a Non-Equilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?" *Econometrica*, 75: 1721-70

[5] Dyer, Douglas, John H. Kagel, and Dan Levin (1989): "A Comparison of Naive and Experienced Bidders in Common Value Offer Auctions: A Laboratory Analysis," *Economic Journal*, 99, 108-115

[6] Eyster, E. and M. Rabin (2005): "Cursed Equilibrium." *Econometrica*, 73(5): 1623-1672.2

[7] Fischbacher, Urs (2007)."z-Tree: Zurich Toolbox for Ready-made Economic Experiments." *Experimental Economics*, 10(2): 171-178.

[8] Hendricks, Kenneth, and Robert H. Porter (1988): "An Empirical Study of an Auction with Asymmetric Information," *American Economic Review*, 78, 865-883

[9] Hendricks, Kenneth, Robert H. Porter, and Bryan Boudreau (1987): "Information, Returns, and Bidding Behavior in OCS Auctions: 1954-1969," *Journal of Industrial Economics*, 35, 517-42

[10] Ivanov, A., D. Levin and J. Peck (forthcoming) "Hindsight, Foresight, and Insight: An Experimental Study of a Small-Market Investment Game with Common and Private Values," *American Economic Review.*

[11] Kagel, John H. (1995): "Auctions: A Survey of Experimental Research," in *The Handbook of Experimental Economics*, ed. by John H. Kagel and Alvin E. Roth. Princeton, NJ: Princeton University Press.

[12] Kagel, J., R. Harstad and D. Levin, (1987): "Information Impact and Allocation Rules in Auctions with Affiliated Private Valuations: An Experimental Study," *Econometrica*, Vol. 55, 1275-1304.

[13] Kagel, John H. and Dan Levin (1986): "The Winner's Curse and Public Information in Common Value Auctions," *American Economic Review*, 76, 894-920.

[14] Kagel, John H., and Dan Levin (2002): "Common Value Auctions and the Winner's Curse." Princeton, NJ: Princeton University Press.

[15] Kagel, John H. and Dan Levin, (forthcoming) "Auctions: A Survey of Experimental Research, 1995 - 2008," in *The Handbook of Experimental Economics, Volume II*, J. H. Kagel and A. E. Roth (eds). Princeton: Princeton University Press.

[16] Laffont, Jean-Jacques (1997): "Game Theory and Empirical Economics: The Case of Auction Data," *European Economic Review*, 41, 1-35

[17] Lind, Barry, and Charles R. Plott (1991): "The Winner's Curse: Experiments with Buyers and with Sellers," *American Economic Review*, 81, 335–46.

[18] McAfee, R. Preston, and John McMillan (1987): "Auctions and Bidding," *Journal of Economic Literature*, 25, 699-738.

[19] McKelvey, R. D. and T. R. Palfrey (1995): "Quantal Response Equilibria for Normal Form Games." *Games and Economic Behavior* 10: 6-38.

[20] Milgrom, Paul and Robert Weber (1982): "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50, 1089-1122.

[21] Nagel, R. (1995): "Unraveling in Guessing Games: An Experimental Study." *American Economic Review*, 85: 1313-1326.

[22] Pesendorfer, Wolfgang and Swinkels, Jeroen M. (1997): "The Loser's Curse and Information Aggregation in Common Value Auctions," *Econometrica*, 65, pp. 1247-1281.

[23] Pevnitskaya, Svetlana (2008): "Effect of Beliefs and Risk Preferences on Bidding in Auctions," Florida State University Working Paper

[24] Stahl, D. and P. Wilson (1995): "On Players' Models of Other Players: Theory and Experimental Evidence." *Games and Economic Behavior* 10: 218-254.

[25] Thaler, Richard H. (1988): "The Winner's Curse," *Journal of Economic Perspectives*, 2, 191-202.

[26] Wilson, Robert (1992): "Strategic Analysis of Auctions," in *Handbook of Game Theory*, Vol. 1, ed. by Robert Aumann and Sergiu Hart. Amsterdam: North-Holland/Elsevier Science Publishers.

7 Appendix: Proofs

Proof of proposition 1

First round of deletion of weakly dominated bid functions: Under the second-price rule, for any x_i, any bid strictly below x_i is weakly dominated (by bidding x_i) since one could lose the auction at a price below x_i even though $x^{max} \geq x_i$. Therefore, we can delete all bid functions, such that $b_i(x_i) < x_i$ for some x_i.

Second round of deletion of weakly dominated strategies: Suppose that bidder i with signal x_i considers bidding $b^+ > x_i$. In the event that bidding x_i wins, bidding b^+ rather than x_i doesn't matter. In the event that bidding b^+ doesn't win, bidding b^+ rather than x_i also doesn't matter.

Now consider the third possible event: that bidding x_i doesn't win but bidding b^+ does. Then, bidder i pays the highest bid among the other $n-1$ players, $\widehat{b}$, where $\widehat{b} \geq x^{max}$. The inequality holds because $\widehat{b} \geq x_i$ (otherwise x_i would have won) and because none of the other bidders ever underbid (by the first round of deletion of weakly dominated bid functions). But then i would make non-positive profits by bidding b^+ whereas she would make zero profits by bidding x_i. Moreover, if $\widehat{b}$ is strictly above x_{max}, then b^+ makes strictly negative profit. Therefore b^+ is weakly dominated and we can delete all bid functions, such that $b_i(x_i) > x_i$ for some x_i.

Proof of proposition 2

A strategy for player i is a probability measure H on $[0,10] \times [0,\infty)$ with marginal cdf on the first coordinate equal to $F(\cdot)$. A pure strategy is a bid function $b : [0,10] \mapsto [0,\infty)$, such that $H(\{x, b(x)\}_{x\ \in [0,10]}) = 1$. That $b(x) = x$ is a NE, follows directly from proposition 1. Here, we prove uniqueness among all symmetric NE.[47]

Assume that H is a symmetric NE. Let $L = \{(x,b)|x \in [0,10], b < x\}$ and $U = \{(x,b)|x \in [0,10], b > x\}$. That is L and U are the sets in $[0,10] \times [0,10]$ strictly below and strictly above the 45^o line, respectively. We need to show that $H(L \cup U) = 0$, or equivalently that $H(L) = 0$ and $H(U) = 0$.

[47]Of course, any bid function which differs from $b(x) = x$ only on a set of measure zero will also be a symmetric NE.

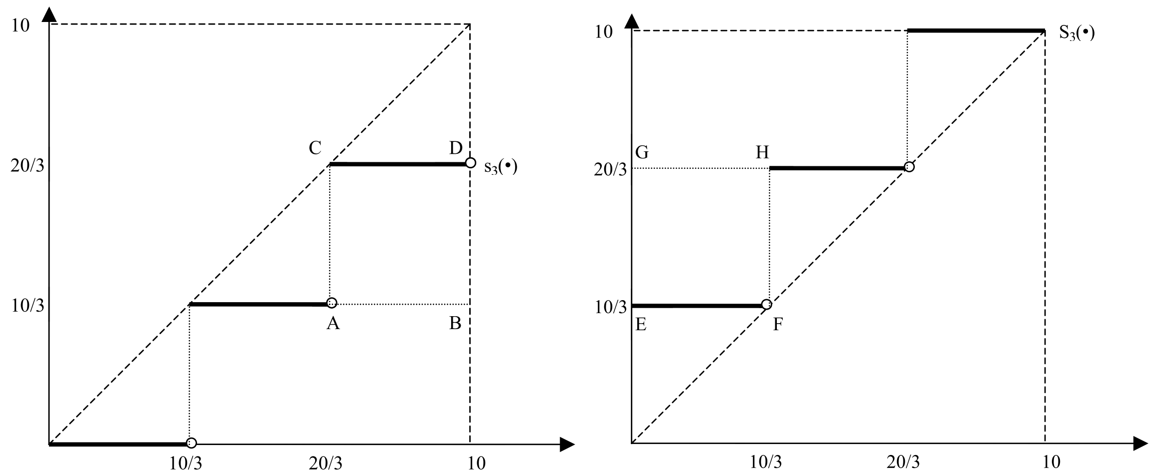

Figure 3:

First, assume $H(L) > 0$. Let $s_k(\cdot)$ be the step function, defined by $s_k(x) = \frac{10}{k} int(\frac{kx}{10})$, where $int(\cdot)$ gives the integer part of a real number ($s_3(\cdot)$ is depicted in the left graph in figure 3). Let $A_k = \{(x,b)|b \leq s_k(x)\} \cap L$, i.e. A_k is the area in L below the $s_k(\cdot)$ function. Note that $k' < k''$ implies $A_{2^{k'}} \subset A_{2^{k''}}$ and that $L = \bigcup_{k \geq 1} A_{2^k}$. Therefore, $H(L) = lim_{k \to \infty} H(A_{2^k}) > 0$.[48] Therefore, for some $\overline{k}$, $H(A_{2^{\overline{k}}}) > 0$. Because $A_{2^{\overline{k}}}$ consists of finitely many rectangles like $ABCD$ in figure 3 ($ABCD$ includes its boundaries, except for point D), it follows that at least one of these rectangles has positive measure. Assume, without loss of generality, $H(ABCD) > 0$.

We will show that, for a positive measure (wrt $H(\cdot)$) of points $(x,b) \in ABCD$, bidding b given signal x is strictly worse than bidding x because there is a positive probability that one will lose the auction to a bid strictly below x. Let $g(\widetilde{b}) = H(\{(x,b)|(x,b) \in ABCD, b \leq \widetilde{b}\})$. Note that $g(\cdot)$ is a non-decreasing function and that $g(\underline{b}) \geq 0$ and $g(\overline{b}) > g(\underline{b})$, where $\underline{b} = min(\{b|(x,b) \in ABCD\})$ and $\overline{b} = max(\{b|(x,b) \in ABCD\})$.

If $g(\underline{b}) > 0$, then $\{(x,b)|(x,b) \in ABCD, b = \underline{b}\}$ has positive measure. For any point (x,b) in this set, bidding $\underline{b}$ given signal x is strictly worse than bidding x since there is a positive probability of a tie at $\underline{b}$.

Assume $g(\underline{b}) = 0$. If $g(\cdot)$ is continuous, choose $b^* \in (\underline{b}, \overline{b})$, such that $0 < g(b^*) <$

[48] To see this, let $B_2 = A_2$ and $B_l = A_l / A_{l-1}$ for $l \geq 3$. Then, $H(L) = H(\bigcup_{l \geq 2} A_l) = H(\bigcup_{l \geq 2} B_l) = \sum_{l \geq 2} H(B_l) = lim_{k \to \infty} \sum_{l=2}^{k} H(B_l) = lim_{k \to \infty} H(A_k) = lim_{k \to \infty} H(A_{2^k})$. The third and fifth equalities follow from the (countable) additivity of probability measures.

$g(\overline{b})$.[49] Then $\{(x,b)|(x,b) \in ABCD, b \leq b^*\}$ and $\{(x,b)|(x,b) \in ABCD, b > b^*\}$ each have positive measure. But then for a positive measure of points (x,b) (the points in the former set), bidding b given signal x is strictly worse than bidding x since there is a positive probability of losing the auction to a bid b, such that $b^* < b < x$.

If $g(\cdot)$ is not continuous, then it has a jump point[50] at, say, b^{**}. Therefore, $\{(x,b)|(x,b) \in ABCD, b = b^{**}\}$ has positive measure. For any point (x,b) in this set, bidding b^{**} given signal x is strictly worse than bidding x since there is a positive probability of a tie at b^{**}. This proves that we cannot have $H(L) > 0$.

The proof that we cannot have $H(U) > 0$ is analogous so that only a brief outline is provided. Assume that $H(U) > 0$. Let $S_k(\cdot)$ be the step function, defined by $S_k(x) = s_k(x + \frac{10}{k})$ ($S_3(\cdot)$ is depicted in the right graph in figure 3). Then, we show analogously to above that a rectangle of the sort $EFGK$ in figure 3, has positive measure. Then, defining $h(\widetilde{b}) = H(\{(x,b)|(x,b) \in EFGH, b \leq \widetilde{b}\})$, we show that for a positive measure (wrt $H(\cdot)$) of points $(x,b) \in EFGH$, bidding b given signal x is strictly worse than bidding x because there is a positive probability that one will win the auction at a price strictly above x^{max}.

Proof of proposition 3

By proposition 1, bidding one's signal is a best-response if others are TL_0's. We now show that, if the other players are TL_0's, then bidding above one's signal is no worse than bidding one's signal.

Suppose that bidder i with signal x_i considers bidding, $b^+ \geq x_i$. In the event that bidding x_i wins, bidding b^+ rather than x_i doesn't matter. In the event that bidding b^+ doesn't win, bidding b^+ rather than x_i also doesn't matter.

Now consider the third possible event: that bidding x_i doesn't win but bidding b^+ does. Then, bidder i pays the highest bid among the other $n-1$ players, $\widehat{b}$, where $\widehat{b} = x^{max}$. The equality follows because (i) others are bidding their signals so that $\widehat{b}$ equals the highest signal among them, and (ii) $\widehat{b} \geq x_i$ (otherwise x_i wins). But then i makes zero profits both by bidding b^+ and by bidding x_i.

[49] This can clearly be done by the intermediate value theorem.

[50] Any nondecreasing function, is either continuous, or has countably many jump points.

Proof of proposition 4

Let $\widehat{B}$ denote the highest bid among the $n-1$ subjects other than i. Given $X_i = x_i$, subject i chooses her bid, b, in order to maximize:

$$\begin{aligned} E(\text{Payoff}|X_i = x_i) &= prob(\widehat{B} < b|X_i = x_i)E(X^{max} - \widehat{B}|X_i = x_i, \widehat{B} < b) \\ &+ \frac{1}{2}prob(\widehat{B} = b|X_i = x_i)E(X^{max} - \widehat{B}|X_i = x_i, \widehat{B} = b) \\ &= prob(\widehat{B} < b)[E(X^{max}|X_i = x_i) - E(\widehat{B}|\widehat{B} < b)] \\ &= \frac{b^{n-1}}{10^{n-1}}[E(X^{max}|X_i = x_i) - \frac{n-1}{n}b] \end{aligned}$$

The second equality follows, because (i) $prob(\widehat{B} = b|X_i = x_i) = 0$, (ii) others' bids (and $\widehat{B}$ in particular) are not informative about X^{max}, and (iii) X_i is not informative about others' bids (and about $\widehat{B}$ in particular). The third equality uses facts about the distribution and expectation of the first-order statistic of $n-1$ *iid* random variables which have the uniform distribution on $[0, 10]$. From the last expression, it is straightforward to verify that the unique optimal bid equals $E(X^{max}|X_i = x_i)$.

8 Appendix: Figures and Tables

	$b < x - 0.25$	$b \sim x$	$x + 0.25 < b \leq 10$	$b > 10$
Part I	17%	23.1%	44.5%	15.4%
Part II	23.3%	22.9%	35.2%	18.6%

Table 5: Percentage of bids in each category for the *ShowBidFn* treatment.

	$b < x - 0.25$	$b \sim x$	$x + 0.25 < b \leq 10$	$b > 10$
Part I	0%	28.3%	60.5%	11.2%
Part II	0%	32.9%	51.1%	16.1%

Table 6: Percentage of bids in each category for the *MinBid* treatment.

Part I / II	*Underbidders*	*Signal-bidders*	*Overbidders*	*Above-10-bidders*	*Indeterminate*	
Underbidders	3	0	0	0	1	**4**
Signal-bidders	0	5	1	0	1	**7**
Overbidders	4	0	10	2	2	**18**
Above-10-bidders	0	1	0	4	0	**5**
Indeterminate	3	1	5	1	2	**12**
	10	**7**	**16**	**7**	**6**	

Table 7: Transition table for the *ShowBidFn* treatment.

Part I / II	*Underbidders*	*Signal-bidders*	*Overbidders*	*Above-10-bidders*	*Indeterminate*	
Underbidders	0	0	0	0	0	**0**
Signal-bidders	0	3	0	1	0	**4**
Overbidders	0	3	14	2	0	**19**
Above-10-bidders	0	1	1	0	0	**2**
Indeterminate	0	1	0	0	0	**1**
	0	**8**	**15**	**3**	**0**	

Table 8: Transition table for the *MinBid* treatment.

Subject ID	% of bids in part II which are a best-response	foregone expected profits (in ECU) in part II	% of $x + 0.25 < b \leq 10$ from part I reduced in part II	% of $x + 0.25 < b \leq 10$ from part I increased in part II
128	27%	3.00	50%	17%
131	0%	10.41	29%	57%
132	0%	13.64	0%	78%
140	18%	8.05	0%	86%
141	18%	2.82	67%	22%
146	0%	6.32	0%	100%
148	9%	475.36	75%	13%
149	0%	10.50	50%	50%
152	27%	1.91	71%	29%
157	55%	2.77	60%	30%
Mean	**15%**	**53.48**	**40%**	**48%**

Table 9: Best-response behavior in part II for subjects who are *Overbidders* in parts I and II in the *ShowBidFn* treatment.(Subject 148 has large foregone expected profits because she bid above 10 on a couple of occasions both in part I and in part II.)

Subject ID	% of bids in part II which are a best-response	foregone expected profits (in ECU) in part II	% of $x + 0.25 < b \leq 10$ from part I reduced in part II	% of $x + 0.25 < b \leq 10$ from part I increased in part II
89	0%	5.11	50%	0%
92	27%	3.18	14%	57%
95	18%	17.91	0%	100%
96	18%	6.73	22%	22%
97	9%	1.78	33%	33%
98	36%	2.73	50%	33%
99	18%	3.32	57%	14%
100	9%	12.68	0%	44%
104	18%	5.64	14%	29%
106	9%	4.36	0%	86%
107	64%	0.27	90%	0%
112	9%	3.00	50%	10%
113	0%	6.40	0%	100%
114	36%	2.50	33%	50%
Mean	**19%**	**5.40**	**30%**	**41%**

Table 10: Best-response behavior in part II for subjects who are *Overbidders* in parts I and II in the *MinBid* treatment.

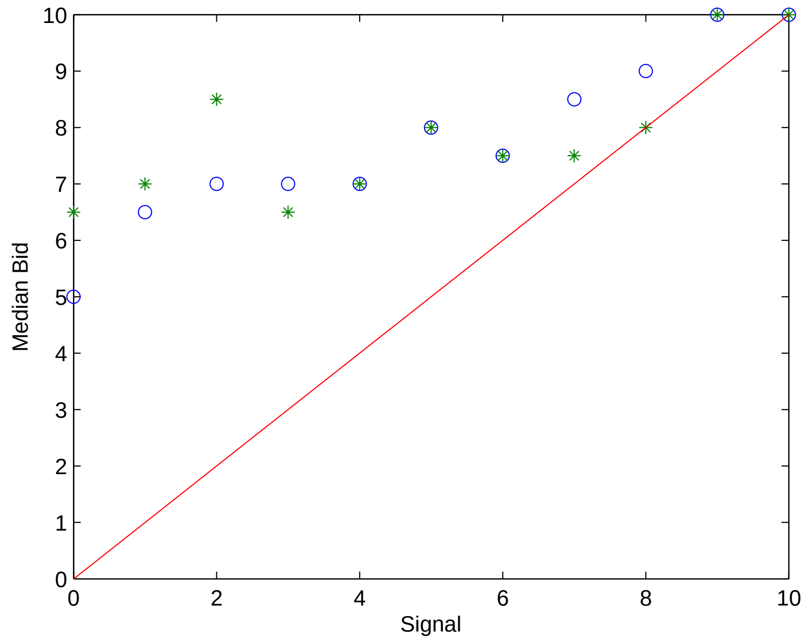

Figure 4: Median bids in parts I (circles) and II (stars) for subjects who are *Overbidders* in parts I and II of the *ShowBidFn* treatment.

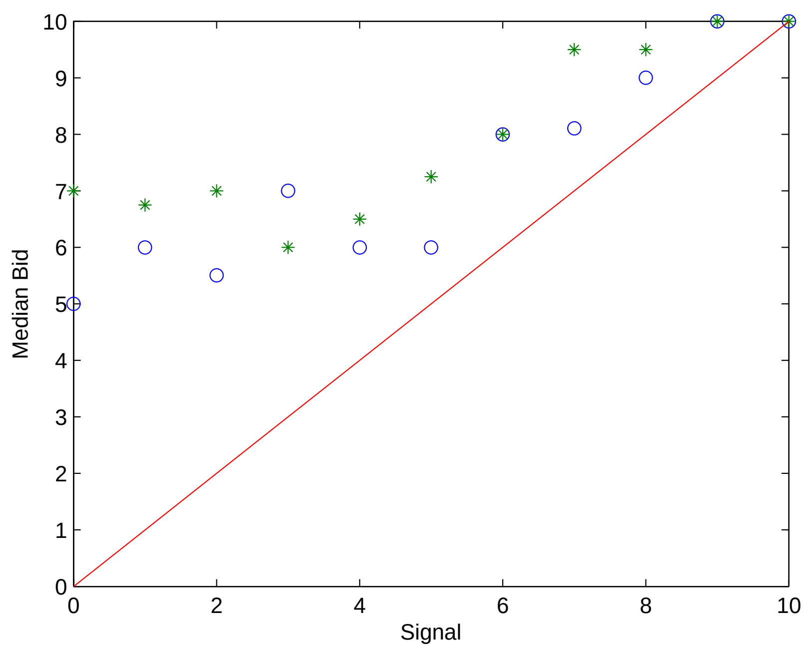

Figure 5: Median bids in parts I (circles) and II (stars) for subjects who are *Overbidders* in parts I and II of the *MinBid* treatment.

9 Appendix: Instructions for *Baseline* treatment

This is an experiment in the economics of market decision-making. The National Science Foundation has provided funds for conducting this research.

This experiment consists of two parts and is expected to last 90 min.

The instructions are simple, and if you follow them carefully and make good decisions, you may earn a CONSIDERABLE AMOUNT OF MONEY which will be paid to you in cash at the end of the experiment.

You will receive a $5 show-up fee which is yours to keep. In addition, you will receive 10 experimental currency units (ECU) starting cash balances for the experiment. You will also have the opportunity to earn ECU in each of the two parts of the experiment. ECU will be converted into dollars at a rate of $0.5 per ECU (i.e. 2 ECU are worth $1). Your total dollar earnings will equal:

$5 show-up fee + 0.5*(10 ECU starting cash balances + ECU earned in part I and part II)

Note that your earnings in part I and part II could be negative (i.e. you could incur a loss) in which case they will be subtracted from your 10 ECU starting cash balances. However, you will receive your $5 show-up fee no matter what.

Caution: This is a serious experiment and talking, looking at others' screens, or exclaiming aloud are not allowed. Should you have any questions please raise your hand and an experimenter will come to you.

Part I

1. In part I of this experiment, we will create a series of auctions in which you will act as bidders for a fictitious item. In each auction, you will be paired randomly with another bidder. A single item will be auctioned off with the two of you as bidders. Your pairings will vary from auction to auction and will remain anonymous.

2. In each auction, you will receive a signal (call it X) and the bidder that you are paired with will also receive a signal (call it Y). X and Y are determined randomly and will lie between 0 and 10. Each whole number within this interval (i.e. 0,1,2,.9,10) has an equal chance of being drawn. In addition, the value of X has no bearing on the value of Y: no matter what the value of X, each whole number between 0 and 10 is equally likely to be the value of Y. The value of the item that is auctioned (call it V) is determined as THE LARGER of the two signals, X and Y.

Prior to bidding in each auction, you will learn X (but not Y); the bidder you are paired with will learn Y (but not X).

Example 1: Suppose you learn that X=6 and the bidder you are paired with learns that Y=4. Then the value of the item is V = 6 ECU.

Example 2: Suppose you learn that X=1 and the bidder you are paired with learns that Y=9. Then the value of the item is V= 9 ECU.

3. Market organization:

In each auction you will submit a bid for the item. The high bidder gets the item and makes a profit equal to the difference between the value of the item and the second highest bid. That is, for the high bidder:

PROFITS = V - (SECOND HIGHEST BID)

If the difference is negative, it represents a loss.

If you do not make the high bid, you will earn zero profits. In this case, you neither gain nor lose money from bidding on the item.

4. Your earnings for part I of the experiment will equal the sum of the profits you made in each auction in part I. (Because your profits in any auction could be negative, your earnings for part I could also be negative.)

5. Even though the computer will keep track of your earnings in each auction, you will not be given any feedback about the outcome of the individual auctions during the experiment.

6. No one may bid less than 0.00 ECU for the item, and bids must be rounded to two digits after the decimal point. You will have 1 minute to place your bid in each auction.

In case of a tie for the high bid, the winner is chosen randomly (50-50 chance). The price the winner pays will be the second highest bid (which is the same as the high bid in case of a tie).

Let us summarize the main points:

1. High bidder gets the item and earns: V - SECOND HIGHEST BID.

2. The value of the item V equals THE LARGER of two signals, X and Y. You learn X; the bidder you are paired with learns Y. The signals are randomly and independently drawn from $\{0, 1, 2, 9, 10\}$.

3. Your earnings for part I equal the sum of the profits you made in each auction in part I. Any questions?

PRACTICE QUIZ

Suppose that Chris and Pat are paired for a given auction. Suppose that Chris has signal X = 7 and Pat has signal Y = 4. Suppose that Chris bids some number B_{Chris} and Pat bids some number B_{Pat}. In addition, suppose for now that Chris' bid is higher than Pat's bid (i.e. $B_{Chris} > B_{Pat}$). Then:

1. The item is obtained by: a) Chris b) Pat

2. The value of the item is V =

3. The second highest bid is: a) B_{Chris} b) B_{Pat}

4. Chris' Profits are: a) 7-B_{Chris} b) 4-B_{Chris} c) 7-B_{Pat} d) 4-B_{Pat} e) 0

5. Pat's Profits are: a) 7-B_{Chris} b) 4-B_{Chris} c) 7-B_{Pat} d) 4-B_{Pat} e) 0

Now suppose that Pat's bid is higher than Chris' bid (i.e. $B_{Pat} > B_{Chris}$). Then:

6. The item is obtained by: a) Chris b) Pat

7. The value of the item is V =

8. The second highest bid is: a) B_{Chris} b) B_{Pat}

9. Chris' Profits are: a) 7-B_{Chris} b) 4-B_{Chris} c) 7-B_{Pat} d) 4-B_{Pat} e) 0

10. Pat's Profits are: a) 7-B_{Chris} b) 4-B_{Chris} c) 7-B_{Pat} d) 4-B_{Pat} e) 0

PART II

1. We will again create a series of auctions in which a fictitious item is sold. Just like in part I, the value of the item (V) in each auction is determined as THE LARGER of two signals, X and Y. X and Y are determined randomly and will lie between 0 and 10. Each whole number within this interval (i.e. 0,1,2,.9,10) has an equal chance of being drawn. In addition, the value of X has no bearing on the value of Y: no matter what the value of X, each whole number between 0 and 10 is equally likely to be the value of Y.

2. In this part of the experiment, instead of bidding for the item against another person, you will be bidding for the item against the computer. The computer will bid by mimicking your bidding behavior from part I of the experiment (as explained below).

Prior to bidding in each auction, you will observe one of the two signals (X). The computer 'observes' the other signal (Y). Then it checks how you bid in part I when you observed that same signal and it makes the same bid.

Example 3: Suppose that the computer observes Y=4. Then the computer checks your bid in part I when you observed signal equal to 4 and it makes the same bid.

3. Market organization:

In each auction, you and the computer will each submit a bid for the item. If you are the high bidder you get the item and make a profit equal to the difference between the value of the item and the second highest bid. That is, if you are the high bidder, you earn:

PROFITS = V - (SECOND HIGHEST BID)

If the difference is negative, it represents a loss.

If you do not make the high bid, you will earn zero profits. In this case, you neither gain nor lose money from bidding on the item.

4. Your earnings for part II of the experiment will equal the sum of the profits you made in each auction in part II. (Because your profits in any auction could be negative, your earnings for part II could also be negative.)

5. Even though the computer will keep track of your earnings in each auction, you will not be given any feedback about the outcome of the individual auctions during the experiment.

6. No one may bid less than 0.00 ECU for the item, and bids must be rounded to two digits after the decimal point. You will have 1 minute to place your bid in each auction.

In case you tie with the computer for the high bid, you win the item with 50% chance. If you win the item, you pay a price equal to the second highest bid (which is the same as the high bid in case of a tie).

We can summarize by saying that the rules for part II are similar to those for part I. The difference is that now you are bidding not against another participant but against the computer which mimics your bidding behavior from part I.

Any questions?

10 Appendix: Screenshot from Part II of the *Show-BidFn* Treatment