

WORKING PAPERS IN ECONOMICS

No. 09/08

EIRIK N. CHRISTENSEN

CURSED EQUILIBRIUM REVISITED



Department of Economics

UNIVERSITY OF BERGEN

Cursed Equilibrium Revisited*

Eirik N. Christensen[†]

Department of Economics, University of Bergen

August 4, 2008

Abstract

Empirical as well as experimental evidence strongly suggests that bidders in common value auctions typically do not conform to the requirements of perfect rationality. Eyster and Rabin (2005) develop a theory and an equilibrium concept — χ -cursed equilibrium — for bounded rational bidding in common value auctions (among other situations), and also present some experimental evidence supporting the theory. This paper comments on these findings of an experiment conducted at the University of Bergen. In the experiment, participants often demonstrate behaviour that is beyond the bounds set by the χ -cursed equilibrium theory, and I present an alternative theory that better explains the experimental findings.

JEL classification numbers: C91, D44.

Keywords: common value auctions, winner's curse, bounded rationality, experiments.

*Without implicating any of them, I would like to thank my colleagues at the University of Bergen and participants at the Gorman Workshop at University of Oxford for valuable comments and suggestions. A very special thanks goes to Steinar Vagstad for helping me to write a first draft of the paper and to Sigve Tjøtta for valuable supervision of my cand.polit. thesis, from which this paper was born. Financial support from the Meltzer Foundation and SNF AS is gratefully acknowledged.

[†]Department of Economics, University of Bergen, Fosswinckels Gate 14, 5007 Bergen, Norway (eirik.christensen@econ.uib.no)

1 Introduction

Economic theory often places rather tough rationality requirements on economic agents, and the theory of common value auctions is a case in point. Rational behaviour in even the simplest form of a common value auction — with a given number of risk-neutral bidders (see e.g., Wilson (1977)) — involves understanding a nexus of intricate effects. One example is the existence of adverse selection: in any symmetric bidding equilibrium, the winner is the one who most overestimates the true value. A rational bidder would therefore bid less than his/her estimate of the true value.¹ The requirements of full rationality seem to be too much for people participating in real-life auctions, and it should therefore not come as a surprise that empirical as well as experimental studies have found many instances of less than perfectly rational behaviour. In controlled experiments with inexperienced bidders in particular, the winner's curse turns out to be more the rule than the exception (See e.g., Kagel, Levin, Battalio and Meyer (1989)).

When theory and reality do not match, the theory will eventually have to give in. One problem facing scholars trying to rebuild their theories to fit better with available data is that while there is in some sense only one way to be rational, behaviour can be boundedly rational in so many ways. An interesting attempt to mend the theory of common value auctions is made by Eyster and Rabin (2005). They allow their bidders to be boundedly rational in the following way: the least rational of their bidders is able to predict the distribution of the other bidders' bids (signals) but does not see any connection between these bids and the underlying information about the true value. In contrast, their most rational players are perfectly rational: they are able both to predict the distribution of bids (signals) and to infer the underlying information, making them conform to the Nash equilibrium. To make their analysis tractable, Eyster and Rabin assume that in any given auction, all bidders are equally rational, somewhere in between these two extremes. They show

¹Another example is the competition effect of more bidders: In private value auctions, if the number of bidders increases, competition flares and each bidder bids more aggressively, *ceteris paribus*. In common value auctions, however, the competition effect may well be dominated by the adverse selection effect: more bidders also mean that the winner — the greatest optimist — will be further from the average expectation.

that the equilibrium bids can be expressed as a weighted average of the bids from the two extremes, with relative weight χ assigned to the least rational alternative. Intuitively, bidders with $\chi > 0$ will be prone to suffer the winner’s curse, and χ is therefore to be interpreted as a “cursedness” parameter. For this reason, Eyster and Rabin let “ χ -cursed equilibrium” denote the equilibrium corresponding to an auction in which all bidders are of type χ . Using data from other studies (Kagel and Levin (1986)), they also present estimates of χ and find that most of the time, the parameter falls within the admissible interval $\chi \in [0, 1]$.

The present paper adds to the χ -cursed equilibrium literature by conducting an experiment tailored to investigate the χ -cursed equilibrium. The new data demonstrate an interesting phenomenon: the estimated χ often fall outside the admissible interval: values of χ exceeding 1 are frequently observed. One interpretation could of course be that in this experiment participants were just performing badly in the experiment, because of their lack of experience, but there is, however, a more plausible explanation: even in the least rational version of χ -cursed equilibrium ($\chi = 1$), bidders are quite sophisticated: they cannot see the connection between other bidders’ bids and their underlying signals, but they are able to predict perfectly the distribution of their bids, which is, after all, quite an achievement. I therefore suggest yet another version of the theory, in which fully rational bidders still play the Nash equilibrium, while the least rational bidders simply bid their estimate. Actual bidders are still assumed to be equally rational and somewhere in between the new boundaries. Actual bids are then still a weighted average of the bids from the extremes, with relative weight ϕ assigned to the least rational alternative. This equilibrium is dubbed ϕ -cursed equilibrium, and the estimates of ϕ typically fall within the admissible interval $\phi \in [0, 1]$.

The paper is organized as follows. In the next section, I briefly present Eyster and Rabin’s (2005) model, including the results that are of importance to the present paper. In Section 3, I present the experiment and demonstrate that i) for a substantial proportion of bidders, the estimated value of χ exceeds 1, suggesting that these bidders are less rational (or less sophisticated) than is possible within the χ -cursed equilibrium framework; ii) there is large variation in the estimated indi-

vidual cursedness parameters within each auction; and iii) the average high bidders have a substantially higher χ than the average bidder. In Section 4, I present the alternative bidding model and the associated ϕ -cursed equilibrium, and demonstrate that most bidders behave according to this theory. Section 5 concludes.

2 A Common Value Auction and the Winner's Curse²

In the following, I will consider a first-price sealed-bid common value auction, in which n risk-neutral bidders bid on an object of value s . Initially every bidder shares the same common prior about s , which, for simplicity, is assumed to be uniformly distributed on the real line \mathbb{R} .³ Before bidding, each bidder i receives a signal x_i that is independently drawn from a uniform distribution over the interval $[s - \frac{a}{2}, s + \frac{a}{2}]$, where a can be interpreted as the underlying uncertainty about the true object value s . Bidders submit bids b_i , the bidder with the highest bid is declared the winner and pays his/her bid to receive the object, while the remaining bidders pay and receive nothing.

There are different ways to define the winner's curse. In particular, because I am dealing with situations in which the common value and the individual signals are random variables, I distinguish between losses that can be attributed to "bad luck" on the one hand and losses that stem from systematic errors made by the bidders on the other. A precise definition is found in Kagel and Levin (2002), which again is based on Capen, Clapp and Campbell's (1971) original idea.

Definition 1 *In any symmetric equilibrium, the winning bidder is the one with the highest signal. Consequently, while the signal is an unbiased estimator for the true value for the average bidder, the winner's signal is an upwardly biased estimator.*

²The following description follows closely that of Eyster and Rabin (2002).

³Strictly speaking, the uniform distribution over the real line is not defined but can be thought of as the limit of the uniform distribution on $[-K, K]$ as $K \rightarrow \infty$, see Klemperer (1999). For practical reasons, in the experiment to be described in the next section, attention is restricted to uniform distributions over a subset of the real line.

The winner's curse refers to the systematic failure to account for this adverse selection effect.

With this definition, a bidder can experience the winner's curse even if his/her winning bid yields positive profit, as long as the average profit is smaller than he/she (erroneously) expected at the time of bidding. The above definition of the winner's curse is often not very useful for practical purposes, because it is very demanding of information. Moreover, in cases in which one has the necessary data, because the winner's curse is associated with less than perfect rationality, the curse might be expected to occur most of the time (or all the time), and sometimes in situations involving substantial gains for the winners. By contrast, the perhaps more intuitive definition of Eyster and Rabin (2005) is tied to whether or not the winner's expected profit is negative.

Definition 2 *Bidder i suffers the winner's curse in the auction's equilibrium if:*

$$E \left[(s - b_i(x_i)) 1_{\{b_i(x_i) > \max_{j \neq i} b_j(x_j)\}} \right] < 0,$$

where $b_i(x_i)$ is the equilibrium bidding function and $1_{\{A\}}$ is the indicator function that takes the value one when A occurs and zero otherwise.

This definition implies that a bidder suffers the winner's curse if the expected value of the object conditional on winning is less than the price conditional on winning.

To analyse the winner's curse, I follow Eyster and Rabin (2005) in constructing what they call a “ χ -virtual” game: the perceived utility of bidder i from winning the auction at a price p when the value of the object is s is given by:

$$(1 - \chi) s + \chi E[s|x_i] - p. \tag{1}$$

This means that bidder i 's perceived valuation of the object is a χ -weighted average of the object's true value and the bidder's expectation of its value given his/her signal. Completely rational bidders have $\chi = 0$ and choose the bid function to maximize the expected value of the object and the winning bid, taking into consideration

the fact that when one's own bid is the winning bid, all other players' bids must be lower and therefore convey information that one's own signal is biased (the highest). By contrast, naïve bidders ($\chi = 1$) choose their bid function to maximize the difference between the expected value of the object and the winning bid, but they do not consider the bad news from having the highest bid, and therefore base their bids on the unconditional expected value of the item. This allows for situations in which the bidders are somewhat naïve, with χ taking any value in the interval $[0, 1]$. However, it is assumed that all bidders are equally rational: different bidders do not have different χ .

In a first-price sealed-bid auction, bidder i chooses his/her bid b_i to maximize:

$$\int_{x_L}^{b_j^{-1}(b_i)} [(1 - \chi) v^n(x_i, y) + \chi r(x_i) - b_i] f_n(y|x_i) dy, \quad (2)$$

where n is the number of bidders, $v^n(x_i, y) = E[s|x_i, \max_{j \neq i}\{x_j\} = y]$ is bidder i 's expectation of the value of the object conditional on his/her signal being x_i and the highest of the other bidders' signals being y , $r(x_i) = E[s|x_i] = x_i$ is bidder i 's expectation of the value of the object conditional on his/her signal x_i , $b_j(\cdot)$ is the common equilibrium bidding function of bidders $j \neq i$ and f_n is the density of y conditional on x_i . With s being drawn from a uniform distribution on \mathbb{R} and the signals subsequently being independent draws from a uniform distribution on $[s - \frac{a}{2}, s + \frac{a}{2}]$, I derive the following equilibrium bid function⁴:

$$b^n(x_i) = x_i - \frac{a}{2} + \chi a \frac{n-2}{2n}. \quad (3)$$

Intuitively, when $\chi = 0$, $b^n(x_i) = x_i - \frac{a}{2}$, which is approximately the risk-neutral Nash equilibrium (RNNE) (see e.g., Milgrom and Weber (1982) and Kagel and Levin (1986)). At the other extreme, if $\chi = 1$, bidders assume that there is no connection between each bidder's action and his/her type (he/she bids because he/she is in a situation with positive private affiliated values). This is the naïve strategic discounting case described by Kagel and Levin (1986), involving $b^n(x_i) = x_i - \frac{a}{n}$. In

⁴Details are found in Appendix A

the latter case, as a response to an increase in the number of bidders, the bidders will bid more aggressively because the competition is increasing. In the former case, the increased competition effect is exactly offset by the adverse selection effect. For a given object value s , the expected highest signal is given by:

$$E[y|s] = s - \frac{(n-1)a}{(n+1)2}. \quad (4)$$

The seller's expected revenue is increasing in n , and his/her expected revenue will be given by:

$$E[b^n(y|s)] = s - \frac{a}{n+1} + a\chi \frac{n-2}{2n}. \quad (5)$$

Therefore, when:

$$n \geq \bar{n} \equiv \frac{2 + \chi + \sqrt{9\chi^2 + 4\chi + 4}}{2\chi}, \quad (6)$$

the seller's expected revenue is larger than s and bidders are facing the winner's curse.⁵ From (3), one may also compute χ by:

$$\chi = \frac{b_i - x_i + \frac{a}{2}}{a \binom{n-2}{2n}}. \quad (7)$$

While Eyster and Rabin (2005) mainly focus on second-price sealed-bid auctions in their paper, they also present data from Kagel and Levin's (1986) experiment to comment on the χ -cursed equilibrium theory for first-price sealed-bid auctions.⁶ Eyster and Rabin find that in 12 of 15 auctions, χ falls inside the admissible interval $\chi \in [0, 1]$, and χ does not seem to be sensitive to changes in the number of bidders.

In Kagel and Levin's (1986) experiment, only 71% of the auctions were won by the high signal holder. This is inconsistent with both RNNE and χ -cursed equilibrium, because these are both symmetric equilibria. It may stem from bidders making errors in bidding (that is, noise), or if bidders have different χ . Hence,

⁵In a symmetric χ -cursed equilibrium, bidders do not fully account for the fact that they only win the object if they have the most positive information of the object. Clearly, any positive χ will lead to the Winner's Curse described in Definition 1, while the winner's curse in Definition 2 will occur iff χ and the number of bidders are high enough.

⁶In the working paper version of "Cursed Equilibrium", first-price sealed-bid auctions are studied more closely (Eyster and Rabin (2002)).

from other experiments, it remains ambiguous whether χ depends on the number of bidders. Kagel and Levin's (1986) results indicated that χ is insensitive to changes in the number of bidders. In the other experiments cited in Eyster and Rabin (2002), uncertainty, a , also varied. To study these results further, the experiment in this paper tests whether χ is dependent on the number of bidders, and hence uncertainty was held constant throughout the experiment.

3 The Experiment

The experiment was conducted with 32 undergraduate economic students at the University of Bergen. The experiment was fully computerized using the software z-Tree (Fischbacher (1999)). No communication between bidders was allowed. The auction structure followed Kagel and Levin (1986), and the experiment lasted for about 1.5 hours. Each bidder was given a paper with information and space for his/her own notes. A short version of the information was also displayed on bidders' computer screens during the experiment. Questions concerning the structure of the experiment were asked and answered privately (on the computer). Before the experiment started, the bidders competed in three 'dry' runs that had no influence on their final payoff (to ensure that they understood the structure and the way the computer program worked). Bidders then answered some control questions before the experiment started. First, 16 of the students (dubbed Session 1) participated in 20 auctions with $n = 4$ (dubbed Series 1) before they bid 20 rounds with $n = 8$ (Series 2). The remaining 16 students (Session 2) did the reverse. Table 1 summarizes the experimental design.

Each auction was a first-price sealed-bid common value auction. In each auction, one item (exemplified as a jar of pennies) was up for sale, and the bidders learned that the true value, s , of this item was drawn from a uniform distribution on the interval $[x_L, x_H] = [50, 250]$ (Norwegian Kroner (NOK)). Next, each bidder i drew independent signals x_i from a uniform distribution on the interval $[s - \frac{a}{2}, s + \frac{a}{2}]$ where a was set to equal 50 NOK in all the auctions.⁷ Finally, the bidders posted

⁷This structure ensures that x_i is an unbiased estimate of the true value s , or can be used to

Session	Series	Auction #	# of bidders		# of observations
			(# of groups)	Experience	(# of winning observations)
1	1	1-20	4 (4)	None	320 (80)
1	2	21-40	8 (2)	Series 1 (20 auctions)	320 (40)
2	3	41-60	8 (2)	None	320 (40)
2	4	61-80	4 (4)	Series 3 (20 auctions)	320 (80)
Total					1280 (240)

Table 1: Experimental design

sealed bids, and the highest bidder won the object and paid his/her bid, and others paid zero and received nothing.⁸

Compared with the auction described in the previous section, in which bidders only know that s is drawn from a uniform distribution over the real line, bidders in the experiment have more precise information. In particular, if they receive a signal that is near the ends of the support for s (that is, if they draw a signal $x_i \in [50, 75]$ or $x_i \in [225, 250]$), the support of the conditional distribution shrinks. Perhaps less intuitive at first glance, also for $x_i \in [75, 225]$, equilibrium bidding is different from what is found in the previous section, because of the fact that bids are interlinked across signals (the optimal bidding functions are necessarily continuous). It can be shown, however, that the equilibrium derived in the previous section is a good approximation of equilibrium bidding in the experimental setting for $x_i > x_L + \frac{a}{2} = 75$ (see Kagel and Levin (1986) for details). I therefore use the bidding functions derived in the previous section to calculate the value of χ for each individual bidder in each auction.

find an unbiased estimate with the boundary values x_L, x_H : Given x_i, a and the boundary values, each bidder could find an upper limit $\min \{x_i + \frac{a}{2}, x_H\}$ and a lower limit $\max \{x_i - \frac{a}{2}, x_L\}$ for the true value s . The limits associated with a given x_i were reported on the screen together with x_i . The distribution of the signal values and the interval $[x_L, x_H] = [50, 250]$ and uncertainty $a = 50$ remained constant during the experiment and were posted as common knowledge. Before each auction, the bidders were informed about the number of bidders in the auction. Bids had to be non-negative and had to be given in whole NOK.

⁸If two or more bidders had the same winning bid, they shared the profit or the loss.

To cover for the possibility of losses, each bidder was given an initial endowment of NOK 200 (approximately \$32 or €25). Profits and losses were added and subtracted from this endowment. If a bidder's endowment became negative, he/she had to leave the experiment. Each bidder's endowment after the experiment was paid as a fee for participating in the experiment.

In each auction round, the following information was displayed on each participant's computer screen *before* bidding: a brief description of the rules of the auction, the number of bidders, own (remaining) endowment, the individual's signal for the present auction, and the highest and lowest possible value of the object corresponding to this signal. *After* bidding, the following information was displayed: the highest (winning) bid, the true value of the object, one's own profit or loss, one's own new endowment, and a history box showing own bids, highest bids, object value and own profit from previous auctions. With this information at hand, each bidder could easily compute the winner's profit, but the identity of the winner remained secret. After the last auction, all bidders answered a questionnaire about how they formed their strategies and how changes in the number of bidders in an auction affected their bidding.

4 Experimental Results

The bidders answered the control questions before the experiment started. In the first session, 11 answered all the questions correctly, and in Session 2, 13 answered correctly. For the bidders with incorrect answers, the error was not of a nature that one would expect to influence their performance. Among the eight bidders who answered incorrectly, six earned more than the average, and the last two earned below average but well above the poorest result. The bidders earned a total of NOK 4824, which gives an average payment of NOK 150.75 to each bidder. None of the bidders went bankrupt.

Auctions with signals above 225 or below 75 are not reported in the results because the bid function described in the previous section is only a good approximation for signals in between these two values, and the bid function therefore cannot be

used to calculate χ . When removing these data, there are 176 winning observations (64 are left out), and 964 bids (316 left out). Tables 2 and 3 summarize the experimental outcomes.

Series (session)	# auctions	Object value average	Signal average (Highest signal average)	Bid average (Highest bid average)	Discounted bid average	Average profit for winner in each auction
1 (1)	46	139.13	137.91 (142.72)	130.98 (143.3)	-6.93	-4.17
2 (1)	34	151.29	150.15 (158.71)	144.06 (159.47)	-6.19	-8.18
3 (2)	31	174.00	172.93 (181.19)	168.76 (183.58)	-3.65	-9.58
4 (2)	65	153.59	152.22 (157.32)	147.43 (156.14)	-4.78	-2.55
Average	176	153.00*	151.73* (166.34)	146.24* (158.26)	-5.42*	-5.30*

*weighted by number of auctions in each series.

Table 2: Summary statistics I

Series (sessions)	# auctions (# bidders)	# auctions with $\pi > 0$ (%)	# auctions won by the high signal holder (%)	# bid $> (4)$ (%)	# auctions with highest bid $> (4)$ (%)	Average π in each auction (t-value)	Average π predicted by RNNE (st.dev)	Average π predicted by strategic discounting (st.dev)	χ (st.dev)
1 (1)	46 (4)	13/46 (28.26%)	22/46 (47.83%)	129/184 (70.11%)	46/46 (100%)	-4.17 (3.21)**	21.41 (4.57)	8.91 (4.57)	2.05 (0.62)
2 (1)	34 (8)	4/34 (11.76%)	21/34 (61.76%)	268/272 (98.53%)	34/34 (100%)	-8.18 (8.49)**	17.59 (2.83)	-1.16 (2.83)	1.37 (0.31)
3 (2)	31 (8)	3/31 (9.68%)	15/31 (48.39%)	237/248 (95.56%)	31/31 (100%)	-9.58 (9.25)**	17.9 (2.52)	-1.04 (2.52)	1.46 (0.28)
4 (2)	65 (4)	28/65 (43.08%)	34/65 (52.31%)	233/260 (89.62%)	65/65 (100%)	-2.55 (3.92)**	21.26 (4.42)	8.76 (4.42)	1.91 (0.4)
Average		48/176 (27.27%)	92/176 (52.27%)	867/964 (89.94%)	176/176 (100%)	-5.3*	20*	5.16*	1.76*

* weighted by number of auctions in each series.

** statistically significant different from zero 0.5% level two-tail t-test.

Table 3: Summary statistics II

Only 48 of the 176 auctions gave the winner a positive profit, and just 52% of the auctions were won by the high signal holder. In all series, the average profit for the winner was negative, with an average loss of NOK 5.30 in each auction. Bidders

did on average bid less than their signal, but the average difference between signal and bid was only NOK 5.42. In this setting, the difference needed to overcome the winner's curse was NOK 15 with four bidders and NOK 20 with eight bidders. . This suggests that the bidders in part ignored the adverse selection problem and faced the winner's curse. Moreover, they did not seem to change their adjustment significantly when the number of bidders was changed. Table 4 presents subjects' profit and their estimated χ for each series.

Session 1					Session 2				
Bidder	Profit	χ series 1 (st.dev)	χ series 2 (st.dev)	χ total (st.dev)	Bidder	Profit	χ series 3 (st.dev)	χ series 4 (st.dev)	χ total (st.dev)
5	209	1.89 (0.10)	1.28 (0.10)	1.55 (0.32)	22	200	0.95 (0.15)	1.06 (0.32)	1.01 (0.26)
8	209	1.27 (0.25)	0.89 (0.12)	1.07 (0.28)	19	198	0.46* (0.35)	1.15 (0.25)	0.83 (0.46)
7	200	0.54* (0.27)	0.36* (0.18)	0.45* (0.24)	32	173	1.30 (0.07)	1.88 (0.10)	1.59 (0.31)
15	200	0.75 (0.32)	0.39* (0.13)	0.55* (0.29)	30	170	1.20 (0.20)	1.84 (0.79)	1.53 (0.66)
14	188	1.85 (0.07)	1.21 (0.05)	1.47 (0.33)	21	169	1.13 (0.33)	1.60 (0.18)	1.36 (0.35)
11	183	1.39 (0.50)	1.17 (0.26)	1.27 (0.39)	18	167	0.99 (0.49)	1.45 (0.61)	1.22 (0.59)
9	173	0.60* (0.55)	0.45* (0.14)	0.51* (0.35)	26	161	1.18 (0.41)	1.76 (0.23)	1.47 (0.44)
10	166	1.10 (0.73)	1.12 (0.43)	1.11 (0.57)	28	161	1.18 (0.39)	2.03 (0.30)	1.61 (0.55)
4	165	1.86 (0.42)	0.94 (0.35)	1.31 (0.59)	31	158	1.29 (0.61)	1.25 (0.34)	1.27 (0.48)
3	148	1.23 (0.98)	1.26 (0.11)	1.24 (0.65)	17	150	1.67 (0.13)	1.90 (0.16)	1.79 (0.19)
6	126	1.35 (1.03)	0.89 (0.39)	1.09 (0.77)	24	138	1.28 (0.47)	1.66 (0.40)	1.47 (0.47)
16	114	1.93 (0.41)	1.44 (0.28)	1.66 (0.42)	20	136	1.47 (0.06)	2.11 (0.12)	1.79 (0.33)
2	102	2.38 (0.13)	1.55 (0.07)	1.89 (0.43)	25	129	0.93 (0.59)	1.04 (0.56)	0.99 (0.56)
12	84	1.27 (0.79)	1.39 (0.21)	1.33 (0.54)	23	121	1.26 (0.40)	1.90 (0.46)	1.59 (0.54)
13	82	1.92 (1.06)	1.16 (0.32)	1.49 (0.81)	27	119	1.39 (0.34)	1.32 (0.57)	1.35 (0.46)
1	31	2.17 (0.76)	1.43 (0.34)	1.76 (0.67)	29	94	1.41 (0.56)	1.77 (0.69)	1.60 (0.65)
Average (st.dev)	148.75 (53.42)	1.47 (0.80)	1.06 (0.44)	1.24 (0.65)	Average (st.dev)	152.75 (28.62)	1.2 (0.46)	1.6 (0.55)	1.40 (0.54)

* Statistically smaller than one at 2.5% significance level one-tail t-test

Table 4: Individual chi and profit in the experiment

4.1 Summary of Findings

From Table 2 - 4 there are several interesting findings.

1. **The experiment participants did not bid in accordance with χ -cursed equilibrium.** Only 3 of the 32 students were found to have their estimated parameter within the admissible interval; for the remaining 29 students, the estimated χ was statistically significantly larger than one. This suggests that typical bidders are more naïve (or less sophisticated) than is possible within the χ -cursed equilibrium framework, which motivates me to extend the model by allowing bidders to be even more naïve.
2. **Bidders seem to be more χ -cursed with four than eight bidders.** This suggests that χ is sensitive to the number of bidders. From the results, it seems that χ declines with the number of bidders (cf. table 4). In the two sessions, 27 of the 32 bidders had a lower χ with eight bidders compared with the setting with four bidders.
3. From Table 4, it is also clear that **there is substantial variation in the individual cursedness parameters within each auction series.** There seems to be an inverse relationship between individual χ and individual aggregate profit: typically, the individuals with a profit of 200 are the least cursed bidders; they do not win auctions in this setting. Moreover, bidders with above-average profit have a lower χ on average than bidders with below-average profit.
4. **Winning bidders typically have a higher χ than losing bidders.** The average high bidder had an estimated χ of 1.76, while the average bidder's χ was 1.33. That only 52% of the auctions were won by the bidder with the highest signal is a further indication of this, because in a symmetric setting, all the auctions should be won by the bidder with highest signal.
5. **Bidding behaviour described by the participants also indicates heterogeneity in χ .** Descriptions from participants in the panel vary from

“bidding as low as possible to earn as much as possible” (earned NOK 114) to “tried to think strategically, bid as much as possible” (earned NOK 31). Bidding patterns varied substantially, and bidders responded to the change in the number of participants in various ways. In the first session, seven bidders responded that the change from four to eight bidders did not affect their bid behaviour. In the second session, only two bidders reported this. In Session 1 (Session 2), eight (five) bidders reported that the increase in the number of bidders probably affected their chance of winning.⁹

In summary: there is heterogeneity among bidders, because χ seems to differ between the individuals. χ -cursed equilibrium seems not to describe the behaviour of the individuals. They fall outside the admissible interval, and χ -cursed equilibrium is not independent of the number of bidders. In the next section, I discuss these findings.

5 A Model for Less Sophisticated Bidders

As shown above, χ -cursed equilibrium does not give a good description of the inexperienced bidders’ behaviour. In particular, most of the participants turn out to be more naïve than what is possible within the χ -cursed equilibrium framework. I here propose an alternative model, based on a different description of the most naïve bidders: they simply bid their signals. Next I follow Eyster and Rabin (2005) in assuming that partly naïve bidders have a bid function that is a weighted average of the RNNE bid function and the bidder’s individual signal x_i .¹⁰ Bidders will then, in a worst-case scenario, not only ignore the adverse selection problem but also ignore the distribution of signals, and will only focus on their own signal as

⁹An interesting bid pattern was described by one of the bidders (earned NOK 173): “In the first rounds, I focused on my signal, and this led to losses. I recognized that there were losses in almost every run, and then decided to play strategically: first I bid NOK 5 over the lowest possible value. Then I changed this to NOK 7 above, then 10 and in the end NOK 12 above the lowest possible value”. This bidder did not change her bidding behaviour when the number of bidders increased. It seems that the bidder was on the way to the Nash Equilibrium, but aggressive bidding from others drove her away from this.

¹⁰True, this may seem a rather extreme description of bidding, but 10% of bids in series 1 and 15% of bids in series 3 were within +/- 1 of the bidders signal.

an unbiased estimator of the true value of the object. The weight assigned to the most naïve bid function will be denoted ϕ , and the corresponding equilibrium will be called a ϕ -cursed equilibrium.

Let $b_{\phi=0}^n = b_{\chi=0}^n = x_i - \frac{a}{2}$ denote the RNNE bid function and let $b_{\phi=1}^n = x_i$ denote the bid function of our most naïve bidders. Then the ϕ -cursed equilibrium bid function is a weighted average of $b_{\phi=0}^n$ and $b_{\phi=1}^n$, the symmetric bid function given by:

$$\begin{aligned} b(x_i) &= (1 - \phi)\left(x_i - \frac{a}{2}\right) + \phi x_i \\ &= x_i - (1 - \phi)\frac{a}{2}, \end{aligned} \tag{8}$$

when $\phi = 0$, the bid function is described by RNNE, and bidding behaviour is independent of the number of bidders (as in the χ -cursed equilibrium). When $\phi = 1$, the bid is independent of the number of bidders and of uncertainty. The bidders only focus on their signal as an unbiased estimator of the true value of the object.

We can now use the bid function in (8) on the experimental data to calculate the values of ϕ for each bidder in each auction, just as when calculating χ . Alternatively one may exploit the functional relationship between the two parameters: because the two bid functions give rise to the same bid in each auction, the following must hold:

$$x_i - \frac{a}{2} + \chi a \frac{n-2}{2n} = b(x_i) = x_i - (1 - \phi)\frac{a}{2} \tag{9}$$

$$\begin{aligned} &\Downarrow \\ \phi &= \frac{n-2}{n}\chi, \end{aligned} \tag{10}$$

that is, ϕ is proportional to χ , where the proportionality factor $\frac{n-2}{n}$ is either $\frac{1}{2}$ (for $n = 4$) or $\frac{3}{4}$ (for $n = 8$). The seller's expected income in a ϕ -cursed equilibrium is given by:

$$E[b^n(y)|s] = s - \frac{a}{n+1} + \phi \frac{a}{2}. \tag{11}$$

The seller’s expected income will then be larger than s and bidders will face the winner’s curse when:

$$n > \frac{2 - \phi}{\phi}, \quad (12)$$

so, when $\phi = 1$, the bidders face the winner’s curse even when there are only two bidders. Because ϕ allows for less-rational bidders than χ , it is clear that it should capture more of the bidder’s behaviour than does χ . This is confirmed in the scatter plot in Figure 1, where more bid data fall within the ϕ -model than within the χ -model.

Table 5 summarizes the estimated ϕ in the experiment, and Table 6 compares the results for winning bidders and all bidders for χ and ϕ .¹¹ By examining these tables, it is apparent that ϕ is less influenced by n than is χ . This is confirmed in our data by a paired t-test; χ is significantly larger with eight bidders compared with four bidders in both Sessions 1 and 2 at 1% significance level. Applying this test to ϕ gives less significant results: the difference for the bidders in Session 1 is insignificant, while in Session 2 it is significant at a 5% significance level.

5.1 Summary of Findings

1. By applying this model, 27 of the 32 bidders fall inside the definition area. **The ϕ -cursed equilibrium seems to give a better description of bidders’ behaviour.** They are far away from that prescribed by the RNNE ($\phi = 0$) and seem to apply a strategy where their own signal is given much weight when they post their bid. Using the connection:

$$\phi = \chi \frac{n - 2}{n},$$

in auctions with four bidders, 28 of the 32 bidders have a ϕ that is smaller than 1 (corresponding to $\chi \leq 2$), and in auctions with eight bidders, 24 bidders have ϕ smaller than 1 ($\chi \leq 1, 33$). This is not surprising given the fact that the ϕ -cursed equilibrium allows bidders to be “less” rational than the χ -cursed

¹¹Note that winning bids are included in the “all bidders average”.





Figure 1: Scatter plot of bid data

Session 1					Session 2				
Bidder	Profit	ϕ series 1 (st.dev)	ϕ series 2 (st.dev)	ϕ total (st.dev)	Bidder	Profit	ϕ series 3 (st.dev)	ϕ series 4 (st.dev)	ϕ total (st.dev)
5	209	0.95* (0.05)	0.96* (0.01)	0.95* (0.04)	22	200	0.71* (0.11)	0.53* (0.16)	0.61* (0.17)
8	209	0.63* (0.13)	0.67* (0.09)	0.65* (0.11)	19	198	0.35* (0.26)	0.57* (0.13)	0.47* (0.23)
7	200	0.27* (0.13)	0.27* (0.13)	0.27* (0.13)	32	173	0.97* (0.05)	0.94* (0.05)	0.96* (0.05)
15	200	0.38* (0.16)	0.29* (0.10)	0.33* (0.13)	30	170	0.90* (0.15)	0.92 (0.40)	0.91 (0.30)
14	188	0.93* (0.03)	0.91* (0.03)	0.91* (0.04)	21	169	0.85* (0.25)	0.80* (0.09)	0.82* (-0.19)
11	183	0.69* (0.25)	0.88 (0.19)	0.80* (0.24)	18	167	0.75* (0.37)	0.72* (0.31)	0.73* (0.33)
9	173	0.30* (0.28)	0.34* (0.10)	0.32* (0.18)	26	161	0.88 (0.31)	0.88* (-0.12)	0.88* (0.23)
10	166	0.55* (0.37)	0.84 (0.32)	0.71* (0.37)	28	161	0.89 (0.29)	1.02 (0.15)	0.95 (0.24)
4	165	0.93 (0.21)	0.70* (0.26)	0.80* (0.26)	31	158	0.97 (0.46)	0.62* (0.17)	0.79* (0.38)
3	148	0.61* (0.49)	0.94* (0.08)	0.80* (0.37)	17	150	1.25 (0.10)	0.95* (0.08)	1.09 (0.18)
6	126	0.67 (0.52)	0.66* (0.29)	0.67* (0.40)	24	138	0.96 (0.35)	0.83* (0.20)	0.89 (0.29)
16	114	0.96 (0.20)	1.08 (0.21)	1.03 (0.21)	20	136	1.11 (0.04)	1.05 (0.06)	1.08 (0.06)
2	102	1.19 (0.06)	1.16 (0.05)	1.17 (0.06)	25	129	0.70* (0.44)	0.52* (0.28)	0.60* (0.36)
12	84	0.63* (0.40)	1.04 (0.16)	0.87* (0.35)	23	121	0.94 (0.30)	0.95 (0.23)	0.95 (0.26)
13	82	0.96 (0.53)	0.87 (0.24)	0.91 (0.39)	27	119	1.04 (0.25)	0.66* (0.28)	0.84* (0.33)
1	31	1.09 (0.38)	1.07 (0.25)	1.08 (0.31)	29	94	1.06 (0.42)	0.89 (0.35)	0.97 (0.39)
Average (st.dev)	148.75 (53.42)	0.74* (0.40)	0.80* (0.33)	0.77* (0.27)	Average (st.dev)	152.75 (28.62)	0.90* (0.34)	0.80* (0.27)	0.85* (0.17)

* Statistically smaller than one at 2.5% significance level one-tail t-test

Table 5: Individual phi and profit in the experiment.

Session	Series	Winning bidder	All bidders	Winning bidder	All bidders
		average χ	average χ	average ϕ	average ϕ
1	1	2.05	1.47	1.03	0.74
1	2	1.37	1.06	1.03	0.80
2	3	1.46	1.20	1.10	0.90
2	4	1.91	1.60	0.96	0.80
Total		1,76*	1,33*	1,02*	0,81*

* Weighted by the number of observations

Table 6: chi and phi for winning bidder and all bidders

equilibrium.

2. **A change in the number of bidders does not seem to influence ϕ ,** suggesting that bidders focus on their own signal and not the adverse selection problem they face. Eyster and Rabin (2002) comment that χ seems to be insensitive to changes in the number of bidders using the data of Kagel and Levin (1986). In this experiment, this is not true: χ decreases as the number of bidders increases. This is also confirmed by a paired t-test for individual bidders in each session.

6 Concluding Remarks

Empirical as well as experimental evidence strongly suggests that bidders in common value auctions typically do not conform to the requirements of perfect rationality. Therefore, descriptive models of such auctions should allow for less than perfectly rational bidding. What is less clear, however, is *how* this should be done. The starting point for this paper is one suggested way to model less-than-perfect rationality, Eyster and Rabin’s (2005) notion of χ -cursed equilibrium. They assume that bidders maximize what they call a virtual utility function, which is a weighted average of an assessment of utility based on a naïve belief about the relationship between one’s own signal and the true value on the one hand and an assessment based on the rational belief on the other hand. Eyster and Rabin proceed to demonstrate that the equilibrium bid function then becomes a weighted average (with the same weights) of the corresponding “naïve” and “rational” bid functions.

The participants in the experiment reported in this paper turned out not to fit the χ -cursed equilibrium model very well: the estimated weights were often outside the admissible interval, a finding that is not very surprising, keeping in mind that even the completely naïve version of χ -cursed equilibrium ($\chi = 1$) bidders are assumed to be quite sophisticated: they perfectly predict the distribution of their competitors’ bids. I also find that the estimated cursedness parameters are shown to depend on the number of bidders, a finding that does not fit the theory.

The broader version of the theory, ϕ -cursed equilibrium, proposed in this paper, in which fully rational bidders still play the Nash equilibrium while the least rational bidders possible simply bid their signals, turns out to fit the experimental data better: most of the out-of-range parameters disappear, and the new cursedness parameter turns out to be more stable across the number of bidders.

References

- Capen, E.C., Clapp, R.V., Campbell, W.N., 1971. Competitive Bidding in High-Risk Situations. *Journal of Petroleum Technology* 23, 641–653.
- Eyster, E., Rabin, M., 2002. Cursed Equilibrium. Working Paper E02-320, IBER, UC Berkeley.
- Eyster, E., Rabin, M., 2005. Cursed Equilibrium. *Econometrica* 73, 1623–1672.
- Fischbacher, U., 1999. z-Tree. Toolbox for Readymade Economic Experiments. IEW Working Paper 21, University of Zürich.
- Kagel, J.H., Levin, D., 1986. The Winner's Curse and Public Information in Common Value Auctions. *American Economic Review* 76, 894–920.
- Kagel, J.H., Levin, D., 2002. Common Value Auctions and the Winner's Curse. Princeton University Press.
- Kagel, J.H., Levin, D., Battalio, R., Meyer, D.J., 1989. First-Price Common Value Auctions: Bidder Behavior and the Winner's Curse. *Economic Inquiry* 27, 241–58.
- Klemperer, P., 1999. Auction Theory. A Guide to the Literature. *Journal of Economic Surveys* 13, 227–286.
- Milgrom P.R., Weber, R.J., 1982. A Theory of Auctions and Competitive Bidding. *Econometrica* 50, 1089–1122.
- Wilson, R., 1977. A Bidding Model of Perfect Competition. *The Review of Economic Studies* 44, 511–518.

Appendix

A The Bid Strategy in χ -Cursed Equilibrium

In the first-price sealed-bid auction, bidder i chooses his/her bid b_i to maximize:

$$\int_{x_L}^{b_j^{-1}(b_i)} [(1 - \chi)v^n(x_i, y) + \chi r(x_i) - b_i] f_n(y|x_i) dy, \quad (13)$$

where $v^n(x_i, y) = E[s|x_i, \max_{j \neq i}\{x_j\} = y]$ is bidder i 's expectation of the value of the object conditional on his/her signal being x_i and the highest of the other bidders' signals being y , $r(x_i) = E[s|x_i] = x_i$ is bidder i 's expectation of the value of the object conditional on his/her signal x_i , $b_j(\cdot)$ is the common equilibrium bidding function of bidders $j \neq i$ and f_n is the density of y conditional on x_i .

Maximizing (2) with respect to the bid b_i yields the following first-order condition (after simplification and exploiting the symmetry condition $b_j^{-1}(b_i) = x_i$):

$$\frac{db^n(x_i)}{dx_i} = ((1 - \chi)v^n(x_i, x_i) + \chi r(x_i) - b^n(x_i)) \frac{f_n(x_i|X_i = x_i)}{F_n(x_i|X_i = x_i)}, \quad (14)$$

where F_n is the C.D.F. that corresponds to the density f_n . With s being drawn from a uniform distribution on \mathbb{R} and the signals subsequently being independently drawn from a uniform distribution on $[s - \frac{a}{2}, s + \frac{a}{2}]$:

$$F_n(x_i|x_i) = \int_{x_i - \frac{a}{2}}^{x_i + \frac{a}{2}} \left(\frac{1}{2} + \frac{x_i - s}{a} \right)^{n-1} \frac{1}{a} ds = \frac{1}{n}, \quad (15)$$

and:

$$f_n(x_i|x_i) = \int_{x_i - \frac{a}{2}}^{x_i + \frac{a}{2}} (n - 1) \left(\frac{1}{2} + \frac{x_i - s}{a} \right)^{n-2} \frac{1}{a^2} ds = \frac{1}{a}. \quad (16)$$

Moreover, $v^n(x_i, x_i) = x_i - \frac{a}{2} + \frac{a}{n}$ (see Eyster and Rabin, ((2002))) and $r(x_i) = x_i$.

Therefore, the first-order condition (3) can be written:

$$\frac{db^n(x_i)}{dx_i} = \left(x_i - (1 - \chi)a \frac{n-2}{2n} - b^n(x_i) \right) \frac{n}{a}. \quad (17)$$

This is a linear first-order differential equation with solution:

$$b^n(x_i) = x_i - \frac{a}{2} + \chi a \frac{n-2}{2n}. \quad (18)$$

Intuitively, when $\chi = 0$, $b^n(x_i) = x_i - \frac{a}{2}$, which is the RNNE as described by, for example, Milgrom and Weber (1982). At the other extreme, if $\chi = 1$, we have a situation where the bidders assume that there is no connection between each bidder's action and his/her type. This is the naïve strategic discounting case described by Kagel and Levin (1986), involving $b^n(x_i) = x_i - \frac{a}{n}$. In the latter case, as a response to an increase in the number of bidders, the bidders will bid more because the competition is increasing. In the former case, the increased competition effect is exactly offset by the adverse selection effect.

B Instructions for the experiment

Below is a translation of the information provided to participants in the experiment before the experiment started. The original information was given in Norwegian.

Introduction

You are now going to participate in a series of auctions. The information leaflet consists of three pages and you can consult this leaflet throughout the experiment. All participants receive the same information. No communication is allowed during the experiment. If you have any practical questions about the information or the experiment, raise your hand and you will be assisted. On the back of this leaflet you can make your own notes if necessary.

Auction background

In the experiment you are going to participate in a series of auctions. In total there are 40 auctions. That implies that you are going to make 40 bids. The experiment is divided into two parts, each with 20 auctions. At the start of each auction, you are told how many participants (including yourself) will be taking part

in the auction. In the first part you will face the same opponents in all 20 auctions. In the second part you will get new opponents for the final 20 auctions.

After the 40 auctions you will be asked to fill out a questionnaire with background information and questions relating to the experiment.

Before we start the experiment everybody has to read through these pages. We will also arrange a trial round with three auctions that is without implications for the results, so that you gain an impression of how the experiment will work. You will also be given some control questions.

Execution of the auction

In all auctions a jar of pennies is sold. The glass itself is worthless; it is just the money in the jar that is of value. The person with the highest bid wins the auction. The person who wins gets the money in the jar but has to pay his/her bid.

Example: If the value in the jar is NOK 100 and highest bid is NOK 80, the winner wins NOK 20. However, if the value in the jar is NOK 100 and the highest bid is NOK 120, the winner loses NOK 20.

Exactly how much money there is in the jar is UNKNOWN and is changed for every auction. The value will always be between NOK 50 and NOK 250. It is the same probability for all whole NOK values between NOK 50 and NOK 250. By that it is meant that it is just as likely that the value is NOK 64 as it is 225.

Before each auction you each get an own signal on the value in the jar. The signal could be higher or lower than the true value in the jar. The signal is maximally NOK 25 above the true value and maximally NOK 25 below the true value. This signal on how much is in the jar tells you what is the highest and lowest value of the jar.

Example: On the computer screen you get a signal: NOK 75. The value that actually is in the jar can then vary over all whole NOK values from NOK 50 to NOK 100. The lowest possible value based on this signal is NOK 50, while the highest possible value based on this signal is NOK 100.

Bids have to be in whole NOK.

Example: One may bid NOK 58 but not NOK 58.50

Everybody gets NOK 200 to use in the experiment. The money is split evenly between the two series, so that one gets NOK 100 for the first 20 auctions and then another NOK 100 for the last 20 auctions. This money does not limit how much you may bid; it is used to compute gains or losses for the winner of the auction. Hence, you may each bid whatever NOK amount you wish.

Example: It is the first auction and the signal one gets is NOK 175. The lowest possible value based on this signal is NOK 150 and the highest possible value is NOK 200. One may bid whatever one wishes; you are not limited by only having NOK 100 in endowment. If the highest bid is NOK 175 and it turns out that the true value is NOK 180, the person with the highest bid wins NOK 5. This sum is added to the NOK 100 in endowment, so that person now has NOK 105. If the highest bid instead were NOK 185, that person would have lost NOK 5 and had NOK 95 when the next auction started.

If you lose all your money you will be eliminated from the experiment.

The money you have when all auctions are completed is your own and will be paid after the experiment is over. Only the winner's endowment is changed in each auction. Those whose bid was not highest in an auction see no change in their endowment. (If there is more than one highest bid, these share the gain or loss evenly.)

After each auction you will be told: how much money there was in the jar, what the highest bid was, your eventual gain/loss and your new endowment. In addition an information box appears that shows the history from previous auctions. Here you are told: your bid, the highest bid, the value in the jar and the winner's profit.

After you have read this information a new round begins.

No communication is allowed during the experiment. If you have any questions raise your hand and you will be assisted.

Good luck.

Department of Economics
University of Bergen
Fosswinckels gate 6
N-5007 Bergen, Norway
Phone: +47 55 58 92 00
Telefax: +47 55 58 92 10
<http://www.svf.uib.no/econ>