

Quantal Response Equilibrium and Overbidding in Private-Value Auctions*

Jacob K. Goeree, Charles A. Holt, and Thomas R. Palfrey

Caltech, Division of Humanities and Social Sciences, 228-77, Pasadena, CA 99125

October 1999

This paper reports the results of a private-values auction experiment in which expected costs of deviating from the Nash equilibrium bidding function are asymmetric, with the implication that upward deviations will be more likely in one treatment than in the other. Overbidding is observed in both treatments, but is more prevalent in the treatment where the costs of overbidding are lower. We specify and estimate a noisy (quantal response) model of equilibrium behavior. Estimated noise and risk aversion parameters are highly significant and consistent across treatments. The resulting two-parameter model tracks both the average bids and the distribution of bids remarkably well. Alternative explanations of overbidding are also considered. The estimates of parameters from a nonlinear probability weighting function yield a formulation that is essentially equivalent to risk aversion in this context. A model in which players experience a "joy of winning" provides a reasonable fit of the data but does significantly worse than the risk aversion model.

1. Introduction

The tendency for human subjects to "overbid" in first-price private value auctions is commonly rationalized in terms of risk aversion (James Cox, Vernon Smith, and James Walker, 1988), an explanation that has led to some controversy. Glenn Harrison (1989), for instance, argues that deviations from the Nash equilibrium may well be caused by low monetary incentives since the costs of such deviations are often rather small. In the particular experiments Harrison reviews, he finds that the median expected cost incurred by a deviation from the risk-neutral Nash prediction is less than five cents.

Harrison's analysis stimulated the most intense cross-fire to date in the experimental economics literature. In one of the comments, Dan Friedman (1992, p.1374) notes that an asymmetry would be needed to explain the observed "misbehavior" as an artifact of payoff

* This research was funded in part by the National Science Foundation (SBR-9818683 and SBR-9631627).

function flatness:

"A simple artifactual explanation of overbidding would be available if the loss function were asymmetric.... In this case expected losses from overbidding a given amount (+) would be smaller than losses from comparable underbidding (-), thus creating a tendency to overbid if subjects are equally likely to lose a penny from overbidding as from underbidding.... Unfortunately, the loss function is almost symmetric.... I conclude that Harrison's simple 'payoff metric' considerations are insufficient to explain *overbidding*."

Our experimental design incorporates Friedman's insight by introducing asymmetries in the expected losses for deviations from a common Nash equilibrium in two auction games. In particular, consider the two loss functions in Figure 1, where the loss is the difference between the expected payoff for a given bid and the expected payoff for the optimal bid, b^* , which is the same in both treatments. Note that the thick line is flatter to the right of b^* , whereas the thin line is flatter to the left. One treatment with more downside risk, described in the next section, has deviation losses shown by the thick line for almost all values a bidder can draw. The other treatment with upside risk, which has the same Nash equilibrium, has a pattern of deviation losses shown by the thin line. Friedman's insight is that, with risk neutrality, we would expect to see underbidding in the upside-risk treatment and overbidding in the downside-risk treatment. If we see overbidding in both treatments, this would suggest the presence of an upward bias caused by risk aversion, a "joy of winning," or other systematic factors. With or without an upward bias, the loss asymmetries in Figure 1 suggest that bids in the downside-risk (thick line) treatment will be higher than those in the other treatment.

Recent advances in the analysis of equilibrium with noisy behavior, called Quantal Response Equilibrium (Richard McKelvey and Thomas Palfrey, 1995, 1998; Jacob Goeree and Charles Holt, 1999a,b), make it possible to specify models that allow 1) deviations from payoff-maximizing choices to be sensitive to deviation costs and asymmetries, and 2) generalized preference structures that include risk aversion, inequality aversion, nonlinear probability weighting, or a utility of winning. Thus explanations based on risk aversion or on noisy responses to small payoff differences are not mutually exclusive. Using a specific parametrization of the quantal response equilibrium, standard maximum likelihood techniques can




Figure 1. Expected Loss as a Function of Bid for the Two Treatments

be employed to evaluate alternative (nested) models. This paper applies these techniques to take a fresh look at the issue of risk aversion versus noisy responses to flat payoffs.

The next section presents the two private-value auction games used in the experiment, along with an analysis of expected payoffs in the unique Nash equilibrium. The experiment and data are described in sections 3 and 4, and section 5 presents a formal model of noisy equilibrium behavior that contains the risk-neutral Nash model as a limiting case. We present estimates of risk aversion and error parameters, which are used to generate predicted bid distributions for each value. Alternative explanations (rank-dependent expected utility, joy of winning, etc.) are discussed in sections 6 and 7, and the final section concludes.

2. Experimental Design: Two Discrete First-Price Auctions

Consider a game in which each of two bidders receives a private value for a prize to be auctioned in a first-price, sealed bid auction. In other words, the prize goes to the highest bidder for a price equal to that bidder's own bid. Each bidder's value for the prize is equally likely to be \$0, \$2, \$4, \$6, \$8, or \$11, as determined by a random device. Bidders know their own values, and bids are constrained to be integer dollar amounts, with ties decided by the flip of a

coin.

Table 1. Equilibrium Expected Payoffs for the Low-Values Treatment
(Optimal Bids Are Denoted by an Asterisk *)

	bid = \$0	bid = \$1	bid = \$2	bid = \$3	bid = \$4	bid = \$5	bid = \$6
value = \$0	0.00*	- 0.25	- 0.83	- 1.75	- 3.00	- 4.58	- 6.00
value = \$2	0.17	0.25*	0.00	- 0.58	- 1.50	- 2.75	- 4.00
value = \$4	0.33	0.75	0.83*	0.58	0.00	- 0.92	- 2.00
value = \$6	0.50	1.25	1.67	1.75*	1.50	0.92	0.00
value = \$8	0.67	1.75	2.50	2.92	3.00*	2.75	2.00
value = \$11	0.92	2.50	3.75	4.67	5.25	5.50*	5.00

The relevant Nash equilibrium in this game with incomplete information about others' preferences is the Bayesian Nash equilibrium, which specifies an equilibrium bid for each possible realization of a bidder's value. It is straightforward but tedious to verify that the equilibrium bids are about half of the private value draw: i.e. bids are \$0, \$1, \$2, \$3, \$4, and \$5 for bidders with values of \$0, \$2, \$4, \$6, \$8, and \$11 respectively.¹ To demonstrate this, we used the conjectured equilibrium bidding strategy to compute expected payoffs for each bid/value combination (see Table I). Each of the six values are equally likely, so the distribution of others' bids is "flat," with a probability of 1/6 for bids of \$0, \$1, ..., \$5. Against this bidding strategy, a bid of \$6 will win for sure, which explains the payoffs in the far right column of Table I. The expected payoffs in the other columns are calculated using the bid probabilities and the fact that

¹ The bids would be exactly half of value if the highest value were \$10 instead of \$11, but we had to raise the highest value to eliminate multiple Nash equilibria (see Appendix A).

a tie results in a coin flip.² The maximum expected payoff in each row coincides with the equilibrium bid, as indicated by an asterisk (*). With a value of \$6, the optimal bid of \$3 results in an expected payoff of \$1.75, and the cost of bidding one unit too high is 25 cents while the cost of bidding one unit too low is only 8 cents, as shown by the thin line in Figure 1. This pattern of asymmetric deviation losses is characteristic, with the 8-cent/25-cent asymmetry holding in all cases except the bottom row.

Recall that Table 1 was constructed on the assumption that the equilibrium bids are \$0, \$1, \$2, \$3, \$4, and \$5 for the six values. When this assumed bidding strategy is used to calculate expected payoffs, the resulting row maxima (marked with asterisks) confirm the original equilibrium. However, this does not exclude the possibility of other symmetric Nash equilibria in pure or mixed strategies, a possibility that is ruled out by the uniqueness proof in Appendix A.

Table 2 shows the analogous calculations for the second treatment, with equally likely private values of \$0, \$3, \$5, \$7, \$9, or \$12. Interestingly, this value increase does not alter the equilibrium bids in the unique symmetric equilibrium, as indicated by the location of optimal bids for each value. Even though the equilibrium bids are the same, we expected more of an upward bias in the second treatment. The intuition can be seen by looking at payoff losses associated with deviations from the Nash equilibrium. Consider, for instance, a bidder with value \$7 whose expected payoffs are shown in the fourth row of Table 2. The Nash equilibrium bid, \$3, is the same as for the value \$6 bidder in the low-values treatment (see row 4 in Table 1). However, while in the low-values treatment the cost of bidding \$1 too high is more than the cost of bidding \$1 too low, the opposite is true in the high-values treatment. The relative costs of deviations for the high value treatment are shown by the dark line in Figure 1, except for the slight change in the bottom row, due to the larger increment between the highest value and the second highest value (needed to rule out multiple Nash equilibria). These asymmetries caused us to expect more overbidding for the high-values treatment.

² For example, consider a bidder with a private value of \$6 (in the fourth row) who faces a rival that bids according to the proposed Nash solution. A bid of 0 has a 1/2 chance of winning (decided by a coin flip) if the rival's value, and hence the rival's bid, is zero, which happens with probability 1/6. Therefore, the expected payoff of a zero bid with a value of \$6 equals $1/2 * 1/6 * (\$6 - \$0) = \$0.50$, as shown in the $b = 0$ column. If the bid is raised to \$1, the probability of winning becomes $1/6 + 1/12 = 1/4$, hence, the expected payoff of a \$1 bid is $1/4 * (\$6 - \$1) = \$1.25$.

Table 2. Equilibrium Expected Payoffs for the High-Values Treatment
(Optimal Bids Are Denoted by an Asterisk *)

	bid = \$0	bid = \$1	bid = \$2	bid = \$3	bid = \$4	bid = \$5	bid = \$6
value = \$0	0.00*	- 0.25	- 0.83	- 1.75	- 3.00	- 4.58	- 6.00
value = \$3	0.25	0.50*	0.42	0.00	- 0.75	- 1.83	- 3.00
value = \$5	0.42	1.00	1.25*	1.17	0.75	0.00	- 1.00
value = \$7	0.58	1.50	2.08	2.33*	2.25	1.83	1.00
value = \$9	0.75	2.00	2.92	3.50	3.75*	3.67	3.00
value = \$12	1.00	2.75	4.17	5.25	6.00	6.42*	6.00

3. The Experimental Procedures

The experiment involved eight sessions, four at Caltech and four at Virginia, each with 10 undergraduate student subjects. No subject participated in more than one of the sessions. Upon arrival, participants were seated in visually isolated booths. We began by reading the instructions in Appendix A. Subjects were told that there would be 15 periods of random pairings determined by draws of numbered ping-pong balls from a bucket. At the start of each period, the experimenter went to each subject's desk and determined a prize value by rolling a six-sided die, in such a way that no other subject could observe the outcome. Once this was done, each subject was required to record this value and write down a bid on his or her information and record sheet. Feasible bids were specified to be (non-negative) integer dollar amounts with a maximum allowed bid equal to the maximum possible value (\$11 or \$12). We then collected the record sheets and determined the winner for each matched pair of subjects. On each information and record sheet, we recorded the subject's own earnings and the "other's bid". The sheets were then returned and we proceeded to the next period. On average, sessions lasted approximately one hour, including the instruction period.

Two sessions at each location were conducted under the low-values treatment ($v = 0, 2, 4, 6, 8, 11$), and two sessions at each location were conducted under the high-values treatment ($v = 0, 3, 5, 7, 9, 12$). At the end of the session subject were paid one-half of their accumulated

profits. Earnings from the auctions averaged \$10.70 in the low-values treatment and \$11.60 in the high-values treatment.³ These amounts are substantially less than expected earnings predicted by the risk neutral Nash equilibrium of \$14.20 and \$17.80, respectively.

4. The Data

Figure 2 shows the period-by-period bids for both treatments, averaged over all sessions.⁴ Consider first the low-values treatment (on the left side). When the value is \$0, there is occasional bidding above value, which vanishes by period 7, as shown by the lowest dark line on the left side. With a value of \$2, bids converge to the Nash level of \$1 by period 5. But with a value of \$4, the bids stay above the Nash bid of \$2, as shown by the second lowest dark line on the left. Overbidding is also observed at higher values. For the high-values treatment on the right side of Figure 2, there is overbidding for every value above \$0 (for which overbidding is dominated by a bid of \$0). To facilitate comparisons across treatments, compare the lowest dark lines, the lowest light lines, etc.

It is apparent from Figure 2 that behavior is erratic in the first few periods. After period 5, bids settle down somewhat and there is no longer a significant trend in the average bids.⁵ Therefore, we use data from the final ten periods to graph average bids as a function of value in Figure 3. For comparison, the Nash equilibrium bidding functions are shown below the lines that track actual bid averages (the thin lines will be discussed below). At this level of aggregation, three key features in the data become clear. First, as remarked in the discussion of Figure 2, in both treatments we observe overbidding relative to the risk neutral Nash prediction. As discussed elsewhere in the literature (e.g. Cox *et al.*, 1992), this is consistent with a variety

³ In addition, subjects received a \$6 show-up fee.

⁴ The complete data set is archived at: <http://www.people.virginia.edu/~cah2k>.

⁵ To check this we fitted a crude regression of bids on time, pooled over all values. The time coefficient is insignificant (slope = - 0.013, with a standard error of 0.009). (Using the reciprocal of time, as suggested by Kagel (1995), we estimated the coefficient of $1/t$ to be 0.07, which is also insignificant with a standard error of 0.1.) In addition, we estimated separate equations for each private value, and the time coefficients for these ten equations are all insignificant. Thus it appears that all aggregate learning has occurred in the first five periods.

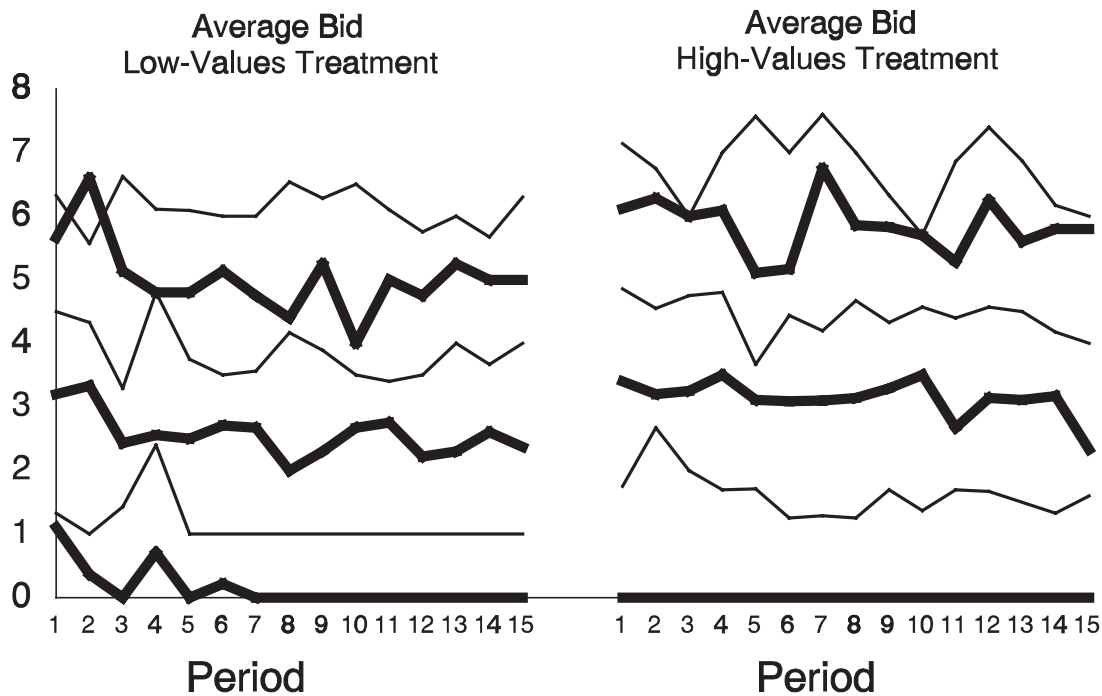


Figure 2. Average Bids by Period

Key: Dark Lines Averages for the First, Third, and Fifth Lowest Values in Each Treatment

of explanations. The most commonly mentioned possibilities are risk aversion and "joy of winning," by which we mean that a player's utility is increased by a fixed amount if they make money in an auction. Second, the extent of overbidding is roughly proportional to value. A pure joy-of-winning model fails to capture this feature, unless somehow the joy of winning is proportional to value for bidders with higher values. On the other hand, with continuous values, constant relative risk aversion results in proportional overbidding and hence is a natural model to consider. Finally, there is relatively more overbidding in the high-values treatment than in the low-values treatment. This feature of the data is consistent with the intuition that noisy behavior will be biased upward in the treatment with more downside risk, an intuition that is captured by the quantal response equilibrium model presented in the next section. These three regularities in the data led us to develop and estimate a hybrid model that combines risk aversion and noisy behavior.

Before proceeding, it is useful to revisit the expected payoff asymmetries identified in Tables 1 and 2. Those tables presented the expected payoff as a function of bid and value, under




Figure 3. Average Bids By Value For Both Treatments.

Key: The QRE/CRRA prediction (thin line) is based on error and risk aversion estimates from the pooled data.

the assumption that opposing bids follow the risk-neutral equilibrium distribution of bids (uniform on $\{0,1,2,3,4,5\}$). As noted above, the data indicate otherwise. This suggests two immediate questions. First, are the equilibrium bid functions still best responses to the empirical bid distributions? Second, is the asymmetry of losses relative to best responses the same as before? The answer to both questions is "yes."⁶ Tables 3 and 4 present the matrix of expected payoffs, analogous to Tables 1 and 2, except for using the empirical bid distribution instead of the equilibrium bid distribution. As before, the results indicate that a risk-neutral bidder would tend to underbid in the low-values treatment, i.e. the expected payoff maxima in each row are at the Nash predictions or below. The reverse is true for the high-values treatment.

⁶ There is only a minor caveat: the optimal response to the empirical bid frequencies in the high-value treatment with a value of \$12 is to bid \$6 instead of the Nash equilibrium bid of \$5, but the difference is only 9-cents.

Table 3. Empirical Expected Payoffs for the Low-Values Treatment
(Optimal Bids Against the Empirical Bid Distribution Are Denoted by an Asterisk *)

	bid = \$0	bid = \$1	bid = \$2	bid = \$3	bid = \$4	bid = \$5	bid = \$6
value = \$0	0.00*	- 0.23	- 0.68	- 1.35	- 2.33	- 3.61	- 5.17
value = \$2	0.15	0.23*	0.00	- 0.45	- 1.17	- 2.16	- 3.44
value = \$4	0.31	0.68*	0.68*	0.45	0.00	- 0.72	- 1.72
value = \$6	0.46	1.14	1.35*	1.35*	1.17	0.72	0.00
value = \$8	0.62	1.59	2.03	2.24	2.33*	2.16	1.72
value = \$11	0.85	2.27	3.05	3.59	4.09	4.33*	4.31

Table 4. Empirical Expected Payoffs for the High-Values Treatment
(Optimal Bids Against the Empirical Bid Distribution Are Denoted by an Asterisk *)

	bid = \$0	bid = \$1	bid = \$2	bid = \$3	bid = \$4	bid = \$5	bid = \$6
value = \$0	0.00*	- 0.20	- 0.58	- 1.27	- 2.22	- 3.44	- 4.88
value = \$3	0.24	0.40*	0.29	0.00	- 0.56	- 1.38	- 2.44
value = \$5	0.41	0.80	0.87*	0.84	0.56	0.00	- 0.81
value = \$7	0.57	1.20	1.45	1.69*	1.67	1.38	0.81
value = \$9	0.73	1.60	2.03	2.54	2.78*	2.75	2.44
value = \$12	0.98	2.20	2.90	3.80	4.45	4.81	4.88*

5. Noisy Decision Making and Risk Aversion

In the high-values treatment, overbidding relative to the Nash equilibrium is consistent with risk aversion but can also be explained by models of "noisy" decision-making. These models capture Friedman's intuition that when bidding too high is relatively more costly than bidding too low, one would expect bids to be biased upwards. Of course, noise in decision making by itself cannot explain the observed overbidding in the low-values treatment, which suggests the importance of other factors such as risk aversion, joy of winning, etc. These alternative explanations do not have to be mutually exclusive and it seems likely that more than

one factor plays a role. In this section we will consider a parametric model that incorporates both noisy behavior and risk aversion. Other possible explanations are considered in the next section.

Following McKelvey and Palfrey (1995), we use a quantal response equilibrium (QRE), in which a probabilistic choice function is combined with a Nash-like equilibrium condition on the consistency of actions and beliefs. A convenient specification is the power-function probabilistic choice rule, which stipulates that the probability of choosing action i , from a set of n possible actions, is proportional to a power function of the associated expected utility, $U^e(i)$:

$$P(i) = \frac{U^e(i)^{1/\mu}}{\sum_{k=1}^n U^e(k)^{1/\mu}}, \quad i = 1, \dots, n. \quad (1)$$

The "error parameter," μ , determines the sensitivity of choice probabilities with respect to payoffs. When $\mu \rightarrow 0$, the option with the highest expected utility is chosen for sure, while in the limiting case of $\mu \rightarrow \infty$, behavior becomes essentially random. For all positive finite values of the error parameter, the choice probabilities in (1) are sensitive to asymmetries of the expected utility function. This insight becomes especially important if one considers the effect of noise in an interactive context such as an auction. For instance, if others' bids are more likely to be upward biased, lower bids have lower expected payoffs and, by (1), are thus chosen less frequently. This makes higher bids even more attractive, which makes them more likely, etc. The quantal response *equilibrium* imposes a consistency condition that the probabilities which determine expected utilities on the right side of (1) match the choice probabilities on the left side that result from probabilistic choice. In some situations, the compounding or "snowball" effects of noise can push the QRE predictions far from Nash.⁷

Following previous literature on auction experiments, we shall consider the specification in which subjects exhibit constant relative risk aversion. With a value ν the expected utility of a bid b is then given by:

⁷ See Capra *et al.* (1999) for a particularly striking example in which the QRE predictions (and data from laboratory experiments) are located at the opposite side of the set of feasible decisions from the Nash predictions.

$$U^e(\mathbf{b}|\nu) = \frac{(\nu - \mathbf{b})^{1-r}}{1-r} P^w(\mathbf{b}), \quad (2)$$

where r is the Arrow-Pratt coefficient of relative risk aversion and the probability of winning, $P^w(\mathbf{b})$, is given by:

$$P^w(\mathbf{b}) = \frac{1}{6} \sum_{\nu \in V} \sum_{\mathbf{b}' < \mathbf{b}} P(\mathbf{b}'|\nu) + \frac{1}{12} \sum_{\nu \in V} P(\mathbf{b}|\nu), \quad (3)$$

where the second term on the right side represents a favorable coin flip in the event of a tie. In our experiments the set of values V is either $\{0,2,4,6,8,11\}$ or $\{0,3,5,7,9,12\}$. We impose an additional condition that bidders do not use bids that could result in negative payoffs.⁸ Given the expected payoffs, the choice probabilities follow from the choice rule in (1), i.e.

$$P(\mathbf{b}|\nu) = \frac{U^e(\mathbf{b}|\nu)^{1/\mu}}{\sum_{\mathbf{b}'=0}^{\nu-1} U^e(\mathbf{b}'|\nu)^{1/\mu}}, \quad \mathbf{b} = 0, 1, \dots, \nu, \quad (4)$$

with the convention that $P(0|0) = 1$.

The QRE can be solved numerically for any values of μ and r , and standard maximum-likelihood techniques can be used to estimate risk aversion and error parameters.⁹ The estimation proceeds by using (4) to calculate decision probabilities for each given parameter combination and then searching for the parameter set that maximizes the product of the probabilities of the decisions observed in the experiment. This simple description somewhat hides the complexity of the estimation, since the equilibrium conditions constitute a collection of simultaneous non-linear equations in the decision probabilities, which enter on both the left and right sides of (4). This system must be solved for each parameter combination.

⁸ In other words, we assume that bidders do not bid above their value. In the last 10 rounds of the data, only 1 out of the 800 bids violates this assumption. Notice that all excluded bids are weakly dominated by bidding 0.

⁹ Numerical computation of equilibrium is feasible only because the strategy space is not too large. This computational feasibility was one factor that led us to choose a coarse grid of values and bids in the experimental design. As mentioned before, we also eliminate dominated strategies.

The maximum likelihood estimates μ and r of are shown in Table 5. That table also reports the value of the likelihood function at these estimates. Asymptotic standard errors are given in parentheses. Notice that the parameter estimates from the two different treatments are of the same order of magnitude even though they are estimated from different data sets. Of particular interest is the fact that the two estimates of r are almost identical, even though there is more overbidding in the high-values treatment. Therefore, we also estimate μ and r using the entire dataset (i.e. forcing the estimates to be the same in both treatments). As expected, this leaves the parameter estimates virtually unchanged.¹⁰ Pooling results in a small (but statistically significant) decline in fit as can be seen by comparing the sum of the loglikelihoods (-752) with the loglikelihood in the pooled case (-760).

The fourth row in Table 5 provides the estimates of the model if we impose the risk neutrality constraint, $r = 0$. The fit is very bad, with a log likelihood of -1138 instead of -760. This reflects the fact that the QRE model with risk-neutral bidders actually predicts *underbidding* in the low-values treatment because the expected payoffs in Table 3 are skewed towards underbidding. In the last row of Table 5, we provide a benchmark to give an idea of how much of the variation in the data is being picked up by the QRE - Risk Aversion model. To construct this benchmark, recall that we are only considering representative-agent models that predict aggregate (not individual) bid frequencies as a function of values. Therefore, the best fit to the data that we could possibly hope for would correspond to a “hindsight” model in which the log likelihood is calculated using the exact frequencies observed in our data. This hindsight model produces a log likelihood of -712. So, the likelihood ratio between the risk neutral model and the hindsight model is -426, while the likelihood ratio of the un-pooled unconstrained QRE - Risk Aversion model and the hindsight model is only -40, a ten-fold improvement.

These parameter estimates can be used to predict average bids for each value, which can then be compared to the empirical averages, to illustrate graphically how good the fit is. The thin lines in Figure 3 give the predicted logit averages for each treatment, where we used the estimates

¹⁰ We also estimated a logit form of the QRE model, in which the power functions are replaced by exponential functions. The risk aversion parameters estimated in this manner were 0.54 for the low-values treatment and 0.56 for the high-values treatment, virtually the same as the estimates in Table 5 with a loglikelihood of -790 for the pooled estimates. We chose to report the estimates from the power function form of QRE only because the fit is slightly better.

Table 5. Maximum Likelihood Estimates for QRE - Risk Aversion Model.
Asymptotic Standard Errors in Parentheses.

	μ	r	Log(L)
Low-values treatment	0.08 (.01)	0.55 (.01)	-333
High-values treatment	0.12 (.01)	0.51 (.02)	-419
Risk Aversion Model (Pooled)	0.10 (.01)	0.52 (.01)	-760
Risk Neutral Model (Pooled)	2.40 (.20)	--	-1138
Hindsight Model (Pooled)	--	--	-712

obtained from the pooled data. The predicted averages are virtually identical to the actual averages, even though the predictions are derived the pooled data.

Not only does the model track conditional bid averages, it also reproduces very well the entire frequency distribution of bidding decisions. We used the power function QRE equations to calculate the complete distribution of bids for each value, as shown in Figure 4. The top part of this figure pertains to the low-values treatment, and the bottom part pertains to the high-values treatment. There are total of 67 bid frequencies; the corresponding predicted probabilities on the right are derived from the two structural parameter estimates.¹¹ When we showed this figure to some of our colleagues, they were unable to guess which side corresponded to the data and which side corresponds to the theoretical predictions. The data are on the left side of the top and bottom panels. While these observations do not rule out other explanations or the possibility that other factors are operating jointly with risk aversion, it is apparent that it will be difficult to find an alternative model that significantly improves the fit obtained with a two-parameter model of risk aversion and noisy behavior.

¹¹ The theoretical predictions in Figure 4 are based on the parameter estimates for the data pooled over both treatments, although cross-treatment predictions are quite similar, as would be expected from the similarity of the parameter estimates.

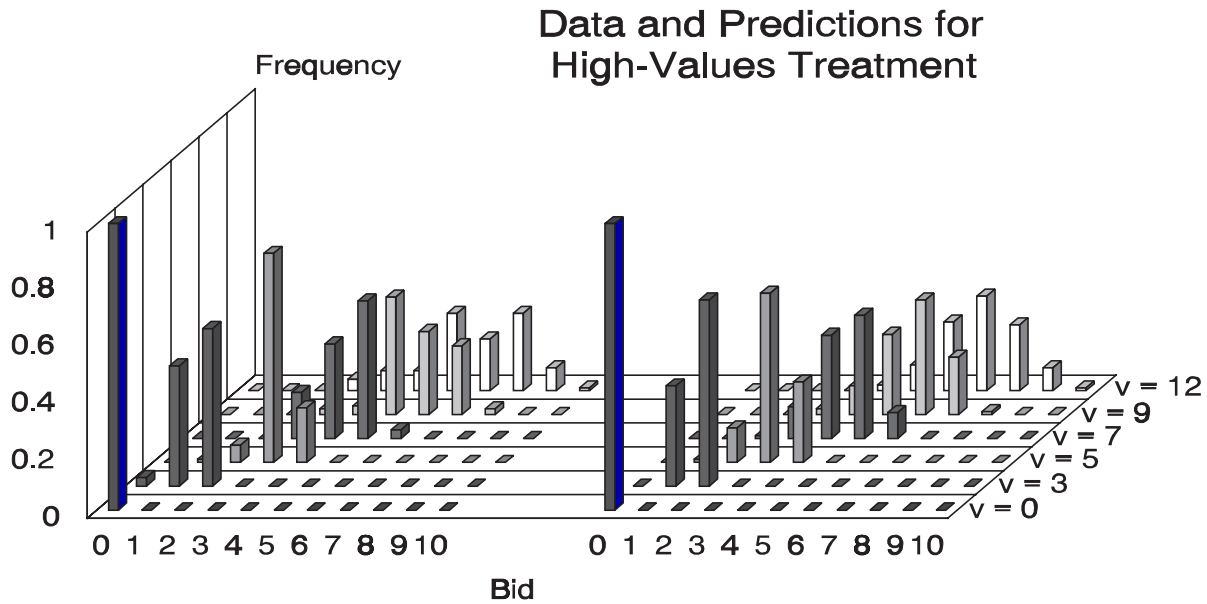
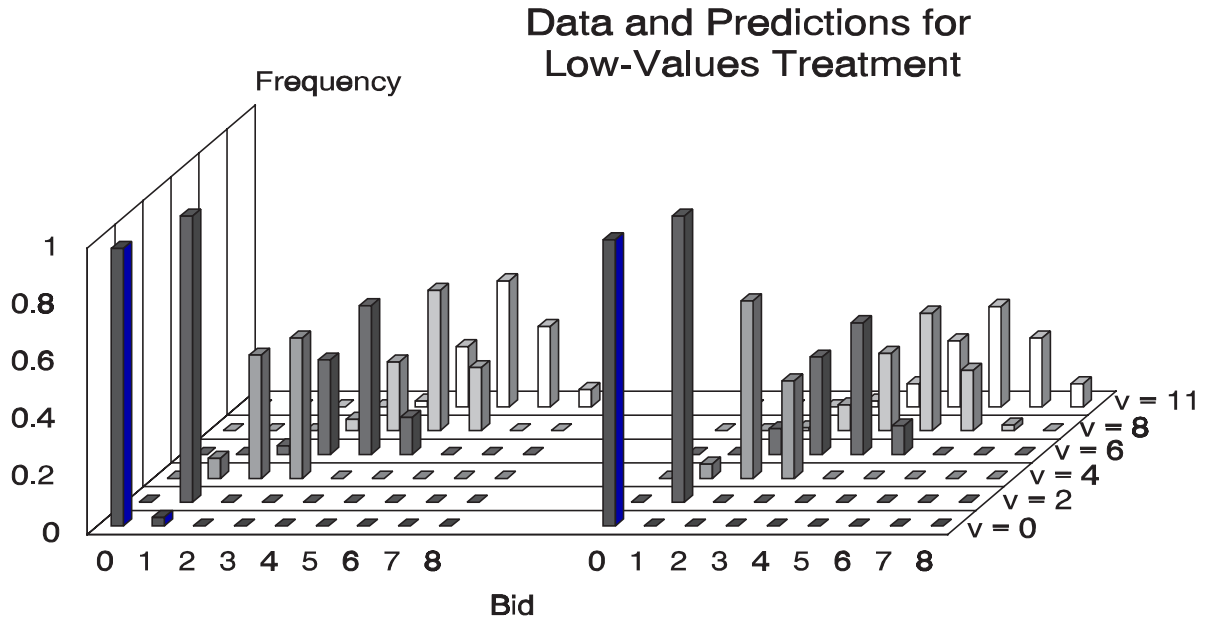


Figure 4. Data (Left) And Predictions (Right) For the Low And the High-Values Treatments.

Recall that we have estimated an equilibrium model with the final 10 periods of data. The equilibrium assumptions, that belief distributions match bid distributions, would be inappropriate if there are discernable time trends in the actual data, which would prevent beliefs from "settling

down." As noted above, there are no time trends after the first five periods. We are not saying that there is no learning in early periods, but rather that the data have stabilized after period 5. The absence of a trend in the average bid data does not necessarily imply an equilibrium situation.

To evaluate this question more carefully, we must examine the bid distributions. Note that an individual bidder sees a private value and considers a bid against the *aggregate* bid distribution, which contains all of the relevant strategic information about others' bids. If an individual's (noisy) best response to the aggregate bid distribution is biased away from that distribution, then the bids would have a tendency to drift in the direction of the best responses. To check this possibility, we compare the actual bid distributions with the quantal response distribution, which by construction, is a noisy best response to itself. The predicted and actual cumulative bid distributions, aggregated across all values, are shown in Figure 5, with the low-values treatment on the left. Using a Kolmogorov-Smirnov test, the null hypothesis that the data are generated by the theoretical distribution, cannot be rejected at the one-percent level.

Of course, adding individual-specific parameters to reflect differences in risk aversion will improve the fit at the individual level, as was found in Cox *et al.* (1983). To verify this, we estimated individual risk aversion parameters, r_i , by calculating expected utilities on the basis of the empirical bid distributions (thick lines in Figure 5), which together with the probabilistic choice rule in (1) determine bid probabilities conditional on μ and r_i . Maximizing the resulting loglikelihood function yields estimates with a median of 0.51, with about three-fourths of the individual estimates in the range from 0.3 to 0.7.^{12,13} While the resulting likelihood function is lower (-574), the eighty-one parameter individual effects model does not improve the fit for the aggregate data.

We share Friedman's uneasiness with specifying a utility function parameter that has the effect of pushing predictions in the direction of the observed data, since "arbitrary special utility

¹² For each range of risk aversion estimates, the number of subjects is shown in parentheses: negative (2 subjects), 0.0-0.1 (3 subjects), 0.1-0.2 (5 subjects), 0.2-0.3 (4 subjects), 0.3-0.4 (14 subjects), 0.4-0.5 (12 subjects), 0.5-0.6 (14 subjects), 0.6-0.7 (15 subjects), 0.7-0.8 (8 subjects), 0.8-0.9 (3 subjects).

¹³ The mean of the individual estimates is .47 with a standard error of .22. Cox and Oaxaca (1996) report an average risk coefficient of .67, using data from the Cox *et al.* (1983) experiment. Chen and Plott (1998) find a mean risk aversion coefficient of .48 for first-price auction experiments with values drawn from a non-uniform distribution.




Figure 5. Aggregate Cumulative Bid Distributions
Key: Data (Thick Line) and QRE Predictions (Thin Line)

effects would ‘explain’ virtually any sort of data” (Friedman, 1982, p.1376). This approach can be particularly *ad hoc* if done without cross checks, as noted by McKelvey and Palfrey (1998) in a different context. It is somewhat reassuring that the risk aversion parameter estimates for the two treatments are virtually identical, even though there is more overbidding in the high-values treatment. This common risk aversion parameter, 0.52, is almost identical to the risk aversion coefficient that Harrison (1990, p.543) concludes would create “virtual equality between observed and predicted behavior” in the four-person experiment from Cox, Roberson, and Smith (1982) that he considers.¹⁴ As noted above, however, our risk aversion estimates are similar to those of Chen and Plott (1998) but somewhat lower than the individual-specific estimates reported by Cox and Oaxaca (1996). Our risk aversion estimates are also essentially the same as those obtained

¹⁴ Harrison notes that an exponent on the utility function of 0.45 reduces the average deviation of the observed and predicted bids to less than a penny. An exponent of 0.45 implies a relative risk aversion coefficient of 0.55, which is close to the estimates in Table 3.

by Holt and Sherman (1997) in a common-value auction experiment.¹⁵ What is apparent from these studies is that bidding behavior is fairly consistent across many auctions with different value structures and different numbers of bidders, and the behavior is consistent with a simple model of risk aversion.¹⁶

Nevertheless, some economists are skeptical that risk aversion would be the primary factor underlying overbidding in low stakes situations, regardless of whether or not they are sympathetic to Harrison's critique. To quote John Kagel and Alvin Roth (1992, p. 1379):

“...risk aversion cannot be the only factor and may well not be the most important factor behind bidding above the risk neutral Nash equilibrium found so often in first-price private value auctions.”

They go on to cite some contradictory evidence. We too are not completely convinced, and therefore we considered some alternative explanations that are discussed in the next two sections.

6. Nonlinear Probability Weighting: An Alternative Explanation?

A major source of skepticism about the risk aversion explanation is that the whole expected-utility framework sometimes provides inaccurate predictions of behavior in simple lottery choice experiments. Several alternatives to expected utility have been proposed. Among these, one of the strongest contenders is *nonlinear probability weighting*, and in particular, the rank-dependent version of nonlinear probability weighting. Any nonlinear weights can be used to

¹⁵ Moreover, Holt and Sherman took subjects who had participated in the common-value auction and put them in an individual decision problem where they received a random private value and made a bid, which would win if it exceeded the realization of a uniform random variable. The optimal bid for a risk neutral person in this context is to bid half of value, but subjects typically bid about 2/3 of value, which is implied by a relative risk aversion parameter of .5. The estimates of individual risk aversion parameters in this individual decision experiment were then regressed with risk aversion parameters estimated for the same people in the common values auction, and the regression coefficient was positive and significant.

¹⁶ Binswanger (1980) provides a direct measurement of risk aversion using high-stakes lottery choice data and concludes that "at high payoff levels, virtually all individuals are moderately risk averse, with little variation according to personal characteristics." All but one of the 118 individuals in his sample of rural Indian farmers exhibited risk aversion, with most common inferred level of relative risk aversion in a "moderate range" (.32 to .81) or an "intermediate range" (.81 to 1.74). Moderate amounts of risk aversion were observed even in low-stakes gambles. Kachelmeier and Shehata (1992) find either risk aversion or risk preference, depending on whether the certainty equivalent of a lottery is elicited in a "willingness to accept" or "a willingness to pay format."

recalculate expected payoffs by replacing the original probabilities over outcomes by transformed probabilities over outcomes. The idea is that individuals may maximize expected utility, but misperceive probabilities. These modified expected utilities are then be inserted into the logit probabilistic choice function to form a stochastic hybrid model, which can be estimated in the same manner as before.

Drazen Prelec (1998) reviews a portion of the existing experimental evidence and concludes that the weighting function $w(p)$ is shaped like an inverted "S": concave for low probabilities, convex for high probabilities, and crosses the diagonal at a probability of about 1/3. Prelec suggests a two-parameter flexible functional form that generates a broad family of curves that includes inverted S-shaped curves:

$$w(p) = \exp(-\beta(-\ln(p))^\alpha), \quad (5)$$

where α and β are parameters of the weighting function. Notice that this function satisfies the endpoint conditions: $w(0) = 0$ and $w(1) = 1$. An inverted S shape will overweight probabilities close to 0 and will underweight probabilities close to 1, which can explain why individuals are willing to bet on low probability gains, and why they shy away from prospects when sure gains are diluted by the small chance of a loss.¹⁷

We estimated the parameters α and β by using this function to transform the win probabilities determined in (3) before using them to calculate expected payoffs, which then

¹⁷ A well-known problem with nonlinear probability weighting, applied directly to the probabilities, is that it does not preserve first-degree dominance relationships, which is the basis of *rank-dependent expected utility*. The preferred solution to this dilemma is to apply the weights to the cumulative distribution function. If we were working with continuous distribution and density functions over money payoffs denoted by $f(x)$ and $F(x)$ respectively, then the probability weighting function approach discussed above is to replace $f(x)$ by $w(f(x))$. In contrast, and the cumulative "rank-dependent" weighting function approach is to replace $F(x)$ by $w(F(x))$, so that probabilities (differences in $F(x)$) are now replaced by $w(F(x))$. In a discrete case, the possible payoffs are ranked in order of preference $x_0 < x_1 < \dots < x_N$, with associated probabilities p_0, p_1, \dots, p_N . As before, we begin with a non-linear weighting function that satisfies $w(0) = 0$ and $w(1) = 1$, as would be the case with (5) for example. But since this function is now applied to the cumulative distribution, we must use differences in the weighting function to obtain transformed probabilities. In particular, the transformed probability for event i is the difference: $p^*_i = w(p_0 + \dots + p_i) - w(p_0 + \dots + p_{i-1})$ for $i = 1, \dots, N$. In a first-price auction, there are only two possible outcomes given a player's value and bid: earnings are either $v - b$ or 0. The fact that the weights sum to 1 implies that the transformed probability of the worst outcome (losing) uniquely determines the transformed probability of the best outcome. Since the transformed probability of losing is multiplied by 0, we are left with only the transformed probability of winning. In other words, ordinary probability weighting and rank-dependent expected utility theory are equivalent in case of first-price auctions.

replaced the expected utility expressions in the numerator and denominator of (4). The likelihood function is constructed in the same manner as before, as a product of the probabilities of the observed bids, which is then maximized over μ , α , and β .

The results of the estimation for the low-values treatment are shown in the top row of Table 6. The weighting function in (5) is quite general and does not require an inverted S shape, and indeed, the estimated parameters produce a function that is essentially quadratic, i.e. convex over the whole range. To see this, note that α and β are essentially 1 and 2 respectively, and that $w(p) = p^2$ in this case. Using this result, the modified expected payoff function becomes: $(v - b)p^2$. Hence, maximizing the modified expected payoff with quadratic probability weighting is equivalent to maximizing untransformed probability times the square root of utility.¹⁸ Put another way, the estimated weight parameters are equivalent to risk aversion, and do not produce the inverted S function that was anticipated.¹⁹

Table 6. Maximum Likelihood Estimates of the Probability Weighting Function. Asymptotic Standard Errors in Parentheses.

	μ	α	β	Log(L)
probability weighting	0.21 (.01)	0.96 (.04)	2.12 (.06)	-759

Thus, the nonlinear probability weighting (or rank dependent expected utility) fits the data just as well as the risk-aversion model. This comes as little surprise since first price auction data cannot identify between nonlinear utility and nonlinear probability weighting. For this reason, we conclude (in the words of Smith and Walker 1993, p. 237) only that bidders behave “*as if risk averse.*” However, it must be emphasized that our estimates indicate a possible problem with the probability weighting model, since the estimates (and even the qualitative shape of the weights)

¹⁸ Note that squaring the utility functions implies that the error parameter is doubled, as can be seen from Table 6.

¹⁹ Other probability weighting functions have been used (e.g. Wu and Gonzalez, 1996; Tversky and Fox, 1994). As a check we also estimated a probability weighting function of the form: $w(p) = ap^b/(ap^b+(1-p)^b)$. The maximum likelihood estimates are: $\mu = .22 (.02)$, $a = .31 (.02)$, $b = 1.46 (.05)$ and $\text{Log(L)} = -759$. A plot of the weighting function that results from these estimates is virtually identical to the simple quadratic function p^2 .

are much different from the parameter values proposed by Prelec, based on data from simple choice experiments. In particular, this raises doubts about the robustness of the inverted S-curve model across different decision environments. Another disadvantage of the QRE - Probability-Weighting model is that it is was a three-parameter model, whereas the risk aversion model fits the data just as well with only two parameters.

7. Models Based on “Joy of Winning”

Another way to rationalize overbidding is if bidders have a jump in their utility function, for positive earnings. This can be interpreted as a *joy of winning*, or more accurately, a joy of “being in the money.”²⁰ This also corresponds to a special form of risk aversion, since the lowest possible earnings a bidder can get (using undominated strategies) is 0, and a joy of winning will shift up the utilities of positive earnings thereby creating a “kinked” utility function that is concave like a risk averse utility function. To measure the degree to which there is an effect of this sort, we incorporate an additional parameter into the QRE - Risk Aversion model. In particular, let expected utility of a bid $b < v$ be given by equation:

$$U^e(b|v) = \left[\frac{(v - b)^{1-r}}{1 - r} + w \right] P(win), \quad (6)$$

where r is the same coefficient of risk aversion as before and $w > 0$ is called the joy of winning.

Table 7 presents the estimates of this three parameter non-linear utility model for the pooled dataset. The coefficient on the joy of winning is low in magnitude and has the wrong sign. The constrained model ($w = 0$) cannot be rejected since the inclusion of w adds little to the likelihood. Moreover, the estimates of μ and r are unaffected by the inclusion of w . Thus the “joy of winning,” at least as we have formulated it here, does not add anything to the explanation of overbidding. However, it is interesting to note that a pure QRE - Joy-of-Winning model (with risk neutrality) actually fits the pooled data quite well, although not as well as the QRE - Risk Aversion model. This is consistent with findings elsewhere. For example, Cox, Smith and Walker (1992) point to an experiment with simulated bidders that also indicates that utility of

²⁰ It can also be interpreted in terms of aspiration levels or threshold utility functions.

Table 7. Maximum Likelihood Estimates for Joy-of-Winning Model.
Asymptotic Standard Errors in Parentheses.

	μ	w	r	Log(L)
Unconstrained	0.10 (.01)	- 0.29 (.07)	0.56 (.01)	-760
Constrained ($w=0$)	0.10 (.01)	--	0.52 (.01)	-760
Constrained ($r=0$)	0.10 (.02)	2.30 (.10)	--	-801

winning models account for bidding behavior less well than risk aversion models.

8. Conclusion

Harrison's "flat maximum" critique is based on the idea that deviations from Nash bidding may not be economically significant if the expected payoff function is flat at the maximum. Friedman added the caveat that a persistent *upward* bias would have to be explained by an asymmetry in the flatness. Here we report an auction experiment that implements such an asymmetry; the two treatments have the same Nash equilibrium bids, but there is more risk for downside deviations in one and more upside risk in the other. Consistent with Friedman's intuition, we observe more overbidding in the downside risk (high-values) treatment, but the absence of underbidding in the other treatment suggests that asymmetric flatness is not the whole story.

Average earnings in the experiment are about 30% below the predictions for a Nash equilibrium with risk neutrality, and there is a large proportional upward bias in bidding for both treatments. Thus a good bidding model for first-price auctions with private values must include some "plus factor" that pushes bids up. Risk aversion is an obvious suspect, given its role in the flat maximum debate and its widespread application elsewhere in economics and finance. A hybrid model, which incorporates constant relative risk aversion into a quantal response equilibrium, produces risk aversion estimates that are virtually the same for the two treatments

(.51 and .55), and are close to those reported in some previous studies. When the parameter estimates for the QRE-risk aversion model are used to predict the distribution of bids conditional on value, the predicted average bids are essentially covered up by the actual data averages in Figure 3. The fit also reproduces the detailed features of the frequency distributions of bids for most of the treatment/value combinations, to such an extent that several of our Caltech colleagues could not distinguish between the actual data and the predicted frequencies in Figure 4.

Given the doubts about whether bidding in other types of auctions is consistent with risk aversion (Kagel and Roth, 1992), we tried two additional approaches. Maximum likelihood estimates for the Prelec (1998) nonlinear weighting function yield a convex shape: the square of the probability. This formulation provides an excellent fit, and the reason is that squared probability times money payoff is just the square of probability times the square root of money payoff, i.e. a constant relative risk aversion of 0.5. These results are promising, with the only concern being that the estimated probability weighting function is quadratic instead of the inverted "S" shape proposed in the literature to explain some anomalies in lottery-choice problems, such as the Allais paradox.

Finally, we tried an alternative model with a "joy of winning," as suggested by Cox, Smith, and Walker (1992). Since this model puts a kink in the utility function, it creates a concavity that has similar qualitative effects on bidding as risk aversion. However, the predicted upward bias is more uniform for all values and does not reproduce the proportional bid increases apparent in Figure 3, so the fit was not as good as for either the risk aversion or probability weighting models.

To summarize, the best explanations of our data involve nonlinearities that add concavity, i.e. bidders behave "as if" they are risk averse. The resulting (two-parameter) model with noisy behavior and some concavity in utility (or convexity in probability) picks up the biases caused by asymmetric payoff flatness. This hybrid model also explains deviations from risk-neutral Nash predictions, both in terms of averages and distributions around those averages.

Appendix A: Nash Equilibrium (for the (0,2,4,6,8,11) treatment)

In this Appendix we derive the unique symmetric Nash equilibrium for the low-values treatment. The proof for the high-values treatment is analogous and is available from the authors on request. We use the following notation:

- Values are indexed by $v \in V = \{0,2,4,6,8,11\}$ with $\text{Prob}(v) = 1/6$ for all $v \in V$
- $P^*(b|v)$ is the equilibrium probability with which a value v bidder selects a bid b
- $P^w(b)$ is the probability of winning with a bid b (which is independent of the value)
- $P(b)$ is the probability that the opponent bids b (so $P^w(b) = \sum_{b' < b} P(b') + P(b)/2$)
- $\pi^e(b|v)$ is the expected payoff with a bid b and value v

The first thing to note is that bidding above value cannot be part of any Nash equilibrium. Suppose, on the contrary, that bids above value could occur in equilibrium. A zero bid will result in a non-negative expected payoff (irrespective of the bidder's value), while bids above value have negative expected payoffs because they have a positive probability of winning (e.g. when the opponent has the same value). Hence, bids above value cannot be optimal. Next, we show that bids are non-decreasing in value: if a bidder with value v bids b with positive probability then bidders with higher values never bid less than b .

Proposition A1. *If $P^*(b|v) > 0$ then $P^*(b'|v') = 0$ for all $b' < b$ and $v' > v$.*

Proof. If $P^*(b|v) > 0$, then a bid equal to b cannot have lower expected payoffs than any other bid, so in particular $\pi^e(b|v) \geq \pi^e(b'|v')$ for all $b' < b$. Since $\pi^e(b|v) = P^w(b) \cdot (v-b)$, we have

$$\frac{P^w(b)}{P^w(b')} \geq \frac{v - b'}{v - b}.$$

Therefore, $\pi^e(b|v) = P^w(b) \cdot (v-b) \geq P^w(b') \cdot (v'-b') \cdot ((v'-b)/(v'-b')) \cdot ((v-b)/(v-b)) > P^w(b') \cdot (v'-b') = \pi^e(b'|v')$, where we used $b > b'$ and $v' > v$ in the next-to-last step. Q.E.D.

Proposition A2. *The unique symmetric Nash equilibrium bids are 0, 1, 2, 3, 4, and 5 for values of 0, 2, 4, 6, 8, and 11 respectively.*

Proof. The method of proof is "bottom-up" and by contradiction: we start with value-2 bidders and show that bids different from 1 are inconsistent with the Nash equilibrium conditions. Then we consider value-4 bidders and show that bids different from 2 lead to a contradiction, etc.

Bidders with value 2 bid 1 in equilibrium.

Expected payoffs for a value-2 bidder are: $\pi^e(0|2) = P(0)$ and $\pi^e(1|2) = P(0) + 1/2 P(1)$. Suppose in contradiction that $P^*(0|2) > 0$, so $P(1)$ has to be 0. But then $P^*(0|2) = 1$ and bidders with values 4 or higher bid at least 2. Payoffs are then $\pi^e(1|4) = 1$ and $\pi^e(2|4) = 2/3 + P(2)$. Hence $P(2) \geq 1/3$, and together with Proposition A1 this implies $P^*(2|4) = P^*(2|6) = 1$. But then $\pi^e(3|6) - \pi^e(2|6) = P(2) - 1/3 + 3/2 P(3)$, which is strictly positive unless $P(2) = 1/3$ and $P(3) = 0$. So bidders with values 8 or 11 should bid at least 4, in which case $\pi^e(3|8) - \pi^e(4|8) = 2/3 - 2 P(4)$, which is strictly positive unless $P(4) = 1/3$, so $P^*(4|8) = P^*(4|11) = 1$. The contradiction finally follows since then $\pi^e(4|11) = 35/6$ while $\pi^e(5|11) = 6$. To conclude, $P^*(0|2) = 0$, or equivalently, $P^*(1|2) = 1$.

Bidders with value 4 bid 2 in equilibrium.

Expected payoffs are: $\pi^e(1|4) = 1/2 + 3/2 P(1)$, $\pi^e(2|4) = 1/3 + 2 P(1) + P(2)$. Furthermore, if $P^*(3|4) > 0$, all bidders with values higher than 4 would bid at least 3, so $P^w(3) < 3/4$ and $\pi^e(3|4) < 3/4$, while $\pi^e(1|4) \geq 3/4$. In equilibrium a bidder with value 4 will therefore bid 1 or 2 (or randomizes between them). Suppose that $p = P^*(1|4) > 0$. If $p < 1$, then $\pi^e(2|4) = \pi^e(1|4)$, which requires $P(2)$ to be $(1 - p)/12$, while $P(2) \geq (1 - p)/6$, a contradiction. Next, if $p = 1$, we must have $1/2 P(1) - 1/6 + P(2) \leq 0$; since $P(1) \geq 1/3$, this implies that $P(1) = 1/3$ and $P(2) = 0$. So bidders with values of 6 and higher bid at least 3. But then $\pi^e(2|6) - \pi^e(3|6) = 1/2 - 3/2 P(3)$, so we must have $P(3) \geq 1/3$, so $P^*(3|6) = P^*(3|8) = 1$. Then $\pi^e(3|8) - \pi^e(4|6) = 1/2 - 3/2 P(3) - 2 P(4)$, which implies $P(3) = 1/3$ and $P(4) = 0$. So a bidder with value 11 bids at least 5, which is contradicted by the expected payoffs of bids of 4 and 5: $\pi^e(4|11) = 35/6$ and $\pi^e(5|11) \leq 33/6$.

Bidders with value 6 bid 3 in equilibrium.

For a value-6 bidder expected payoffs are: $\pi^e(2|6) = 4/3 + 2 P(2)$, $\pi^e(3|6) = 1 + 3 P(2) + 3/2 P(3)$. If $P^*(4|6) > 0$ all bidders with values higher than 6 would bid at least 4, so $P^w(4) < 5/6$ and $\pi^e(4|6) < 5/3$, while $\pi^e(2|6) \geq 5/3$. And a bid of 5 yields at most 1, so, in equilibrium, a bidder

with value 6 bids 2 or 3 (or randomizes between them). Suppose that $q = P^*(2|6) > 0$. If $q < 1$, then $\pi^e(2|6) = \pi^e(3|6)$, which requires $P(3)$ to be $(1 - q)/9$, while $P(3) \geq (1 - q)/6$, a contradiction. Next, if $q = 1$, we must have $P(2) - 1/3 + 3/2 P(3) \leq 0$; since $P(2) \geq 1/3$, this implies that $P(2) = 1/3$ and $P(3) = 0$. So bidders with values of 8 or 11 bid at least 4. But then $\pi^e(3|8) - \pi^e(4|8) = 2/3 - 2P(4)$, so we must have $P(4) \geq 1/3$, so $P^*(4|8) = P^*(4|11) = 1$. But expected payoffs for the value-11 bidder are then: $\pi^e(4|11) = 35/6$ and $\pi^e(5|11) \leq 6$, a contradiction.

Bidders with value 8 bid 4 in equilibrium.

Expected payoffs are: $\pi^e(3|8) = 5/2 + 5/2 P(3)$, $\pi^e(4|8) = 2 + 4P(2) + 2P(4)$. If $P^*(5|8) > 0$ a bidder with value 11 would bid at least 5, so $P^w(5) < 11/12$ and $\pi^e(5|8) < 33/12$, while $\pi^e(3|8) \geq 35/12$. And bids above 5 yield expected payoffs of at most 2, so, in equilibrium, a bidder with value 8 bids 3 or 4 (or randomizes between them). Suppose that $r = P^*(3|8) > 0$. If $r < 1$, then $\pi^e(4|8) = \pi^e(3|8)$, which implies that $3/2 P(3) - 1/2 + 2P(4) = 0$, or $P(4) = (1 - r)/8$ while $P(4) > (1 - r)/6$, a contradiction. If $r = 1$, then $\pi^e(3|8) - \pi^e(4|8) = 1/2 - 3/2 P(3) - 2P(4)$, and since $P(3) \geq 1/3$, this is negative unless $P(3) = 1/3$ and $P(4) = 0$. But this means that a bidder with value 11 bids at least 5, but then $\pi^e(4|11) = 35/6$ while $\pi^e(5|11) \leq 33/6$, a contradiction. So $P^*(4|8) = 1$.

Bidders with value 11 bid 5 in equilibrium.

For a value-11 bidder expected payoffs are: $\pi^e(4|11) = 14/3 + 7/2 P(4)$, $\pi^e(5|11) = 4 + 6P(4) + 3P(5)$. Bids of 6 or higher have expected payoffs less than 5, and since $\pi^e(4|11) \geq 63/12$ they cannot occur in equilibrium. A bidder with value 11 will therefore bid 4 or 5, or mix between them. Suppose that $s = P^*(4|11) > 0$. If $s < 1$, then $\pi^e(4|11) = \pi^e(5|11)$, which implies that $5/2 P(4) - 2/3 + 3P(5) = 0$, or $P(5) = (13 + 5s)/36 > (1 - s)/6$, a contradiction. If $s = 1$, then $\pi^e(4|11) = 35/6$ but $\pi^e(5|11) = 6$, the final contradiction. Q.E.D.

Appendix B: Instructions (for the (0,2,4,6,8,11) treatment)

id: _____

We will begin with a series of auctions, and you will be identified by the ID number shown at the top of the page. In each auction you will be paired with another participant, using draws of numbered ping pong balls. Each ping ball is marked with one of the ID numbers of the ten people in the room. Here is a bucket of ping pong balls, and we will draw them two at a time to determine who is matched with whom. The two people who are matched will bid for a prize, as explained below.

In each auction, you will make a monetary bid for a prize that will be awarded to the highest bidder. Your prize value will be determined by a throw of a 6-sided die. A throw of 1 will determine a value of \$0, a throw of 2 will determine a value of \$2, a throw of 3 will determine a value of \$4, a throw of 4 will determine a value of \$6, a throw of 5 will determine a value of \$8, and a throw of 6 will determine a value of \$11. We will begin by coming to each of your desks to throw the die to determine the value of the prize for each of you. The value of the prize to you will be equally likely to be \$0, \$2, \$4, \$6, \$8, or \$11, which you will know before you make your bid, but you will not know the value of the prize to the person you are matched with. All you know is that the value to them is equally likely to be \$0, \$2, \$4, \$6, \$8, or \$11.

After each person finds out their own prize value, you will choose a bid in an integer dollar amount, which is recorded below. Then we will collect these sheets and use draws of numbered ping pong balls to match you with another bidder. The prize goes to the higher bidder. There are likely to be some ties, since every bid must be in integer dollar amounts (\$0, \$1, ...), and in the event of a tie we will decide who wins with the flip of a coin (heads and the person with the higher ID number wins, tails and the person with the lower ID number wins). If you are the high bidder (or win the flip in the event of a tie), you earn the difference between your own prize value and your bid, and if you are the low bidder, you earn nothing in this auction.

earnings = your own prize value - your bid

(if you are the high bidder or win the flip in case of a tie)

earnings = 0

(if you are the low bidder or lose the flip in case of a tie)

We will conduct 15 auctions, and your cash earnings for this part of the experiment will be determined by adding up the earnings for each auction and dividing by 2. (If you bid above your prize value and win the auction, your earnings will be negative and will be subtracted to determine the total, which is then divided by 2.) You can use the record sheet that is attached to record your prize value and your bid for each auction. After you see your value and choose a bid, we will collect the sheet, match you with another bidder, and return the sheet to you with the other person's bid and your earnings.

Are there any questions? Please do not talk with others during the experiment. This part of the experiment will be followed by another, quite different decision making experiment. And your total cash earnings for all of the decisions made today will be paid to you in cash before you leave the room.

References

- Binswanger, Hans P. (1980) "Attitude Toward Risk: Experimental Measurement in Rural India," *American Journal of Agricultural Economics*, 62, August, 395-407.
- Capra, C. Monica, Jacob K. Goeree, Rosario Gomez, and Charles A. Holt (1999) "Anomalous Behavior in a Traveler's Dilemma?" *American Economic Review*, 89(3), June, 678-690.
- Chen, Kay-Yut and Charles R. Plott (1998) "Nonlinear Behavior in Sealed Bid First-Price Auctions," *Games and Economic Behavior*, 25, 34-78.
- Cox, James C., and Ronald L. Oaxaca (1996) "Is Bidding Behavior Consistent with Bidding Theory for Private Value Auctions," in R. M. Isaac, ed., *Research in Experimental Economics*, Vol. 6, Greenwich, Conn.: JAI Press, 131-148.
- Cox, James C., B. Roberson, and Vernon L. Smith (1982) "Theory and Behavior of Single Price Auctions," in V. L. Smith, ed., *Research in Experimental Economics*, Vol. 2, Greenwich, Conn.: JAI Press, 1-43.
- Cox, James C., Vernon L. Smith, and James M. Walker (1988) "Theory and Individual Behavior of First-Price Auctions," *Journal of Risk and Uncertainty*, 1, March, 61-99.
- Cox, James C., Vernon L. Smith, and James M. Walker (1992) "Theory and Misbehavior of First-Price Auctions: Comment," *American Economic Review*, 82, 1392-1412.
- Davis, Douglas D. and Charles A. Holt (1993) *Experimental Economics*, Princeton, N.J.: Princeton University Press.
- Friedman, Daniel (1992) "Theory and Misbehavior of First-Price Auctions: Comment," *American Economic Review*, 82(5), December, 1374-1378.
- Goeree, Jacob K. and Charles A. Holt (1999a) "Stochastic Game Theory: For Playing Games, Not Just For Doing Theory," *Proceedings of the National Academy of Sciences*, 96, September, 10564-10567.
- Goeree, Jacob K. and Charles A. Holt (1999b) "Asymmetric Inequality Aversion and Noisy Behavior in Alternating-Offer Bargaining Games," forthcoming in the *European Economic Review*.
- Harrison, Glenn W. (1989) "Theory and Misbehavior of First-Price Auctions," *American Economic Review*, 79, 749-762.
- Harrison, Glenn W. (1990) "Risk Attitudes in First-Price Auction Experiments: A Bayesian Analysis," *The Review of Economics and Statistics*, 541-546.

- Holt, Charles A. and Roger Sherman (1997) "Naive Bidding and the Winner's Curse in Auctions with Independent Common Value Components," Discussion Paper presented at the 1995 CREED Conference on Experimental Economics, Amsterdam.
- Kachelmeier, Steven J. and Mohamed Shehata, "Examining Risk Preferences Under High Monetary Incentives: Experimental Evidence from the People's Republic of China, *American Economic Review*, 82(5), December, 1120-1141.
- Kagel, John H. (1995) "Auctions: A Survey of Experimental Research," in J. Kagel and A. Roth (eds.), *Handbook of Experimental Economics*, Princeton: Princeton University Press, 1995, 501-585.
- Kagel, John H. and Alvin E. Roth (1992) "Theory and Misbehavior of First-Price Auctions: Comment," *American Economic Review*, 82, 1379-1391.
- McKelvey, Richard D. and Thomas R. Palfrey (1995) "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10, 6-38.
- McKelvey, Richard D. and Thomas R. Palfrey (1998) "Quantal Response Equilibria in Extensive Form Games," *Experimental Economics*, 1(1), 9-41.
- Merlo, Antonio and Andrew Schotter (1992) "Theory and Misbehavior of First-Price Auctions; Comment," *American Economic Review*, 82(5), December, 1413-1425.
- Ochs, Jack (1994) "Games with Unique, Mixed Strategy Equilibria: An Experimental Study," *Games and Economic Behavior*, 10, 202-217.
- Prelec, Drazen (1998) "The Probability Weighting Function," *Econometrica*, 66(3), May, 497-527.
- Smith, Vernon L. and James M. Walker (1993) "Rewards, Experience, and Decision Costs in First-Price Auctions," *Economic Inquiry*, 31, 237-245.
- Tversky, Amos and C. R. Fox (1994) "Weighting Risk and Uncertainty," *Psychological Review*, 102, 269-283.
- Wu, George and R. Gonzalez (1996) "Curvature of the Probability Weighting Function," *Management Science*, 42, 1667-1690.