NBER WORKING PAPER SERIES

ARE STRUCTURAL ESTIMATES OF AUCTION MODELS REASONABLE? EVIDENCE FROM EXPERIMENTAL DATA

Patrick Bajari
Ali Hortacsu

Working Paper 9889
http://www.nber.org/papers/w9889

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2003

We would like to thank Lanier Benkard, Liran Einav, Joris Pinkse and Peter Reiss for helpful comments. Bajari would like to acknowledge the National Science Foundation, grants SES-0112106 and SES-0122747, for generous research support. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

Are Structural Estimates of Auction Models Reasonable? Evidence from Experimental Data
Patrick Bajari and Ali Hortacsu
NBER Working Paper No. 9889
July 2003
JEL No. L0, C5, D8

ABSTRACT

Recently, economists have developed methods for structural estimation of auction models. Many researchers object to these methods because they find the rationality assumptions used in these models to be implausible. In this paper, we explore whether structural auction models can generate reasonable estimates of bidders' private information. Using bid data from auction experiments, we estimate four alternative structural models of bidding in first-price sealed-bid auctions: 1) risk neutral Bayes-Nash, 2) risk averse Bayes-Nash, 3) a model of learning and 4) a quantal response model of bidding. For each model, we compare the estimated valuations and the valuations assigned to bidders in the experiments. We find that a slight modification of Guerre, Perrigne and Vuong's (2000) procedure for estimating the risk neutral Bayes-Nash model to allow for bidder asymmetries generates quite reasonable estimates of the structural parameters.

Patrick Bajari
Department of Economics
Stanford University
Stanford, CA 94305-6072
and NBER
bajari@econ.duke.edu

Ali Hortacsu
Department of Economics
University of Chicago
1126 East 59th Street
Chicago, IL 60637
and NBER
hortacsu@uchicago.edu

1 Introduction

In modern industrial organization, there has recently been a growth in the number of empirical papers that use a structural econometric modelling approach to analyze data on firm and consumer behavior. One of the most active research areas in this line of work has been on the analysis of auction data. As pointed out by the influential survey of Laffont and Vuong (1996), auction models appear especially well-suited to base structural econometric models on "because of the availability of many data sets and the well-defined game forms associated with auctions."

The literature on structural estimation of auctions began with Paarsch (1992) who estimated parametric models of private and common value first-price auctions. The ensuing literature is too large to cite completely here. However, a sampling of papers that structurally estimate first-price auction models include Donald and Paarsch (1993, 1996), Elyakime, Laffont, Loisel and Vuong (1994), Flambard and Perrigne (2001), Campo (2001), Guerre, Perrigne and Vuong (2000), and Hendricks, Pinkse and Porter (2002). These papers all assume that bidders are playing a Bayes-Nash equilibrium to the auction game, and estimate the distribution of bidders' private information by inverting the equilibrium mapping from private information to observed bids. However, none of these papers are able to compare their estimates of bidders' valuations to the bidder's "true" private information. Therefore, it is not possible to directly assess whether these structural estimates are reasonable.

Despite the growth of structural econometric modelling in empirical analysis of auctions and other industrial organization contexts, many researchers are not comfortable with the strict rationality assumptions imposed in structural estimation. This lack of comfort is not wholly unwarranted. For example, the equilibrium bid function in a first-price auction game is the solution to a complicated differential equation. Even the most ardent supporter of rationality might object to an econometric procedure that assumes that observed behavior corresponds to the solution of such a complicated model.

In this paper, our goal is to assess whether structural models of first-price auctions can generate reasonable estimates of bidders' private information. We structurally estimate four alternative first-price auction models using bid data from the experiments of Dyer, Kagel, and Levin (1989). The four models we estimate are: 1) risk neutral Bayes-Nash, 2) risk averse Bayes-Nash, 3) Quantal Response Equilibrium (QRE), and 4) an adaptive model of learning. In experimental data, unlike field data, we can directly observe bidder valuations. In order to assess which model generates the best estimates of the structural parameters, we measure the distance between the estimated and "true" valuations. We chose to use the Dyer, Kagel, and Levin experiment since the number of bidders varies, which, as we shall discuss in the paper, is required for the identification of some of the models.

We begin by structurally estimating two standard models utilized in the econometric literature: Bayesian-Nash equilibrium with risk neutral bidders and Bayesian-Nash equilibrium with risk averse bidders. We estimate the risk neutral model using nonparametric methods suggested by Elyakime, Laffont, Loisel and Vuong (1994) and Guerre, Perrigne, Vuong (2000). Also, we consider a slight modification to this estimator, similar to the approach used by Pesendorfer and Jofre-Bonet (2003) and Bajari and Ye (2003) that allows for bidder asymmetries. In the asymmetric auction model, bidders are assumed to draw their valuations from independent, but not identical distributions. The source of bidder asymmetries in these experiments is not clear. It could be that some subjects experience a slight "thrill" from winning the auction. Regardless of the source, in the course of conducting this research we discovered (quite by accident!) that modifying the estimator of Guerre, Perrigne and Vuong (2000) in this way leads to better estimates of the structural parameters.

We use the methods of Campo (2001) to estimate the risk averse model. Risk-aversion has been suggested by Cox, Smith and Walker (1983a,b;1985a,b;1988) as an explanation to a large number of first-price experiments in which bids were found to be higher than risk neutral Bayes-Nash bids (this is referred to as the "overbidding" puzzle). Cox, Smith and Walker propose a model with constant relative risk aversion (CRRA) utility and bidder-specific risk aversion coefficients – this is the model estimated by Campo's procedure.

Since many economists are skeptical of the rationality assumptions typically used in structural modeling, we also estimate two models where bidders are less than perfectly rational. The first "less-than-rational" modelling approach we take is one in which bidders do not correctly anticipate the distribution of competitors' bids, as Bayesian-Nash equilibrium posits. Instead, we assume that bidders estimate the probability of winning with a given bid from previous plays of the game and then maximize expected utility subject to these beliefs – a type of "adaptive best-response" strategy. This approach is inspired by Sargent (1993) who suggested models where agents form beliefs as econometricians would form beliefs. Similar approaches have been used in Bray (1982), Marcet and Sargent (1989), Marcet and Nicolini (1997), and Cho and Sargent (1997).

The second "less-than-rational" model we use is the Quantal Response Equilibrium (QRE) model of McKelvey and Palfrey (1998). The QRE model generalizes rational models of behavior in games by allowing a logit error term to influence players' decisions in the spirit of discrete choice models. Following McKelvey and Palfrey's original paper, the QRE model has quickly gained a large degree of acceptance among experimental economists due to its ability to fit a large amount of previously puzzling experimental data through a low-dimensional relaxation of the benchmark Nash equilibrium model. (See Goeree and Holt (1999) and Capra, Gomez, Goeree, and Holt (2002)). Goeree, Holt, and Palfrey (2002) is the only experimental paper we are aware of that applies the QRE framework to an auction setting. They report that a model in which bidders have a common risk aversion coefficient, along with the QRE concept adapted to an incomplete information setting, provides a good explanation to the overbidding phenomenon.[2]

This paper makes two main contributions to the auction literature. First, we find that if we slightly modify the Elyakime, Laffont, Loisel and Vuong (1994) and Guerre, Perrigne and Vuong (2000) estimator of the risk neutral Bayes-Nash model to allow for bidder asymmetries, the estimated distribution of valutions is quite close to the true distribution of valuations. To the best of our knowledge, previous work has not compared the estimated valuations to the true valuations in order to assess whether structural auction models yield reasonable results.

Second, we illustratate some potential strengths and weakness of other structural models of bidding. To the best of our knowledge, previous papers have not evaluated whether "behavioral" models of bidding, such as the QRE or the learning model, generate better structural estimates than rational models of bidding, such as the risk neutral Bayes-Nash or risk averse Bayes-Nash.

Several words of caution are in order when interpreting our results. First, we recognize that experimental environments may differ significantly from "real" economic environments. In real auctions, the stakes

[2] However the crucial difference between Goeree, Holt and Palfrey's (2002) work and ours is that their econometric analysis uses data on bidders' private valuations to assess the fit of the QRE model.

are much higher, and there is a lot more room for heterogeneity and unobserved environmental factors to confound the econometric specification. However, our finding that the benchmark risk neutral Bayesian-Nash equilibrium model performs quite well, even in this experimental setting, is encouraging for present and future users of structural econometric tools.

Second, our exclusive focus on the model's ability to recover structural parameters is limited in scope. Models of behavior in auctions have many different uses beyond structural estimation, including theoretical analysis of bidding, forecasting behavior, and market design. Just because a particular model performs well at one task does not automatically imply that it will perform well at other tasks. As a practical matter, it is wise for an economist to consider the robustness of his results to alternative modelling assumptions, including weakening rationality assumptions.

The plan of the paper is as follows: in sections 2 to 5, we discuss each of the four model specifications in turn, and either review or develop econometric estimators for the structural parameters of interest. In sections 6 and 7, we report our estimates and assess the strengths and weaknesses of all four models. Section 8 concludes.

2 The Bayes-Nash Equilibrium Model

In this section, we describe the standard model of first-price sealed-bid auctions with risk neutral bidders and two methods that can be used to nonparametrically estimate this model. In the model, there are $i = 1, ..., N$ symmetric bidders with valuation v_i for a single and indivisible object. Bidders' valuations are private information that is independently and identically distributed with cdf $F(v)$ and pdf $f(v)$. The support of v_i is the interval $[\underline{v}, \overline{v}]$.

In the auction, bidders simultaneously submit sealed bids b_i. If i's bid is the highest, her utility is $v_i - b_i$, and is 0 otherwise.[3] Bidder i's von Neumann-Morgenstern (vNM) utility is therefore

$$u_i(b_1, ..., b_n, v_i) \equiv \begin{cases} v_i - b_i & \text{if } b_i > b_j \text{ for all } i \neq j \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

Attention is restricted to the symmetric equilibrium of this game where the equilibrium bid function is $b = \mathbf{b}(v)$. Under weak regularity conditions, the equilibrium bid function is strictly increasing and differentiable so that its inverse $\phi(b)$ exists and inherits these properties. Let $\pi_i(b_i; v_i)$ denote the expected profit of bidder i as a function of i's bid and valuation, assuming that all other players bid according to $b = \mathbf{b}(v)$. Then

$$\pi_i(b_i; v_i) \equiv (v_i - b_i)F(\phi(b))^{N-1}. \tag{2}$$

In equation (2), bidder i's expected utility is i's surplus $v_i - b_i$, conditional on winning, times the probability that bidder i wins the auction.

2.1 Structural Estimation

In the Bayes-Nash equilibrium model above, the first order condition for maximizing expected profits (2) implies that

$$v = b + \frac{F(\phi(b))}{f(\phi(b))\phi'(b)(N-1)} . \tag{3}$$

Let $G(b)$ and $g(b)$ be the distribution and density of the bids, respectively. Since $G(b) = F(\phi(b))$ and $g(b) = f(\phi(b))\phi'(b)$ we can rewrite (3) as

$$v = b + \frac{G(b)}{g(b)(N-1)} . \tag{4}$$

[3] If a tie occurs, the object will be awarded at random among the set of high bidders. However, ties have zero probability in equilibrium.

An implication of equation (3), first exploited by Elyakime, Laffont, Loisel and Vuong (1994), and expanded upon in Guerre, Perrigne and Vuong (2000), is that it is possible to use an observed sample of bids to estimate $F(v)$, the distribution of private information. Suppose that the econometrician observed a sample of bids $b_{(1)}, ..., b_{(T)}$ generated from $K = \frac{T}{N}$ repetitions of the auction (where bidders repeatedly draw iid valuations from $F(v)$). Let $\hat{G}(b)$ and $\hat{g}(b)$ denote estimates of $G(b)$ and $g(b)$. For the t^{th} bid, by evaluating equation (4), one can estimate the private information of the bidder associated with the t^{th} bid, denoted as $\widehat{v}_{i,t}$ as follows:

$$\widehat{v}_{i,t} = b_{(t)} + \frac{\widehat{G}(b_{(t)})}{\widehat{g}(b_{(t)})(N-1)}. \tag{5}$$

Given T estimated valuations, $\widehat{v}_{i,1}, ..., \widehat{v}_{i,T}$, the economist then estimates $F(v)$ assuming that the valuations are i.i.d.

The above estimation can be done nonparametrically, using standard kernel methods to estimate $\widehat{G}(b)$ and $\widehat{g}(b)$ from observed bids. Similarly, a kernel method is used to estimate $F(v)$ from $\widehat{v}_{i,1}, ..., \widehat{v}_{i,T}$. The details of our estimation procedure will be described in section 6.

2.2 An Asymmetric Estimator

In the course of conducting this research, we found that a slight generalization of the methods of Guerre, Perrigne and Vuong (2000) can sometimes lead to improved estimates of the structural parameters. The generalization that we propose is similar to the estimation strategy used in Pesendorfer and Jofre-Bonet (2003) and Bajari and Ye (2003) which allows bidders to be asymmetric.

In the asymmetric auction model, bidders' valuations are independently distributed, but we no longer require them to be identically distributed. We let F_i denote the cdf for bidder i's valuation v_i. In the asymmetric model, the distribution of valuations, F_i is common knowledge so that each bidder knows that the distribution of valuations varies across bidders.[4]

Since the valuations are no longer symmetric, it follows that the bidding strategies are no longer symmetric. Let $G_i(b)$ denote the cdf for player i's bid. Then player i's expected utility from bidding b is

$$\pi_i(b_i; v_i) = (v_i - b) \prod_{j \neq i} G_j(b). \tag{6}$$

Note that in (6), unlike in (2), bidder i takes account of the fact that strategies systematically differ across bidders. Heterogeneity is the norm rather than the exception in many economic problems. In a given strategic situation, agents are likely to behave differently due to preferences, previous experience or even their personality. In order for bidder i to behave optimally, she should take account of this heterogeneity when making a best response.

The first order conditions for maximizing utility are

$$v_i = b_i + \frac{1}{\sum_{j \neq i} \frac{g_j(b_i)}{G_j(b_i)}}. \tag{7}$$

Let $\widehat{G}_i(b)$ and $\widehat{g}_i(b)$ denote estimates of the density and distribution of bids. Then, analogously to (5) we estimate $\widehat{v}_{i,t}$ as

$$\widehat{v}_{i,t} = b_{(t)} + \frac{1}{\sum_{j \neq i} \frac{\widehat{g}_j(b_{(t)})}{\widehat{G}_j(b_{(t)})}}. \tag{8}$$

Note that (8) collapses to (5) when bidders are assumed to be symmetric so that $G_i(b) = G(b)$.

It is not obvious why bidders in an auction experiment should be asymmetric. One possibility is that

[4] For a comprehensive review of the theory of the asymmetric auction, the reader can consult Maskin and Riley (2000a,b).

some bidders enjoy participating in the experiments and receive a thrill from winning the auction. This is not a priori implausible given the low stakes in most experiments and the fact that the subjects are typically students. Our motivation for presenting the asymmetric estimator is that, regardless of the reason that bidders are not symmetric in the experiments, equation (8) generates better estimates of the structural parameters than (5).

It is worth noting that in order to estimate the model with asymmetries, we must estimate $G_i(b)$ separately for each bidder i. In the symmetric model, on the other hand, since the distribution of bids is identical for each agent, the number of observations required to precisely estimate $F(v)$ is lower. Therefore, while the asymmetric model is more flexible than the symmetric model, we pay a price: there must be a large number of observations for each agent in the data set. While some empirical data sets have this feature, many do not.

3 Risk Aversion

The second model considered is Bayes-Nash equilibrium with risk averse bidders. A regularity in first-price sealed-bid auction experiments is that the observed bids tend to be higher than the equilibrium bids. Cox, Smith and Walker (1985, 1988) note that one possible explanation for overbidding is risk aversion, and offer a model with heterogeneous risk averse bidders to explain experimental results in first-price auctions.

Since the experimental literature focuses on models with bidder specific constant relative risk aversions, we are going to implement the method suggested by Campo (2001).[5] Following Cox, Walker and Smith (1985, 1988), Campo (2001) assumes that bidder i has a constant-relative-risk-aversion (CRRA) utility function, $U_i(x) = x^{\theta_i}$. In this specification, $1 - \theta_i$ is bidder i's coefficient of relative risk aversion, with

[5] Another notable attempt to structurally estimate a model of first-price IPV auction models with risk aversion is Campo, Guerre, Perrigne, and Vuong (2000) who reject the hypothesis of risk neutral bidders in U.S. Forest Service timber sale auctions.

$\theta_i = 0$ corresponding to risk neutrality. As in Campo (2001), the valuations are assumed to be i.i.d. However, since bidders are asymmetric, the Bayesian-Nash equilibrium is no longer symmetric. As in the previous section, let the cdf of i's bids be $G_i(b)$ and let $g_i(b)$ denote the density. Define

$$Y_i(b) = \sum_{j \neq i} \frac{g_j(b)}{G_j(b)}. \tag{9}$$

Then, the first order condition is

$$v_i = b_i + \frac{\theta_i}{Y_i(b_i)}. \tag{10}$$

Observe that when bidders are risk neutral, that is $\theta_i = 1$ and $G_i = G_j, \forall i, j$, then (10) reduces to (3).

3.1 Structural Estimation

For $\alpha \in [0, 1]$, let $b_{i\alpha}$ and $b_{j\alpha}$ be defined as solutions to the following implicit equations

$$G_i(b_{i\alpha}) = G_j(b_{j\alpha}) = \alpha. \tag{11}$$

i.e. $b_{i\alpha}$ and $b_{j\alpha}$ are the α-th percentiles of the individual bid distributions of i and j. Let v_α be the α-th percentiles of the valuation distribution, satisfying the implicit equation $\alpha = F(v_\alpha)$. Then, assuming that the equilibrium bid functions $b_i(v)$ and $b_j(v)$ are strictly increasing, we get that $G_i(b_{i\alpha}) = G_j(b_{j\alpha}) = \alpha = F(v_\alpha)$. It then follows from the first order conditions for agents i and j, (10), that

$$b_{i\alpha} - b_{j\alpha} = \frac{\theta_j}{Y_j(b_{j\alpha})} - \frac{\theta_i}{Y_i(b_{i\alpha})}. \tag{12}$$

The key insight of Campo (2001), is that equation (12) can be used to estimate θ_i for each bidder. Once θ_i is estimated for each bidder, equation (10) can be used to estimate $F(v)$ as in the risk neutral case. Observe that with N bidders, one can write equation (12) in $\frac{N(N-1)}{2}$ ways for each α. Since there are only N parameters, $\{\theta_i\}_{i=1:N}$ to identify, the relative risk aversion coefficients are over-identified.

Campo (2001) suggests the following procedure to estimate $F(v)$ and $\{\theta_i\}_{i=1:N}$:

1. Estimate bidder specific bid distributions, $\{\widehat{G}_i(b)\}_{i=1:N}$. This can be done nonparametrically using kernel methods.
2. Fix α, and find $\{\hat{b}_{1\alpha}, ..., \hat{b}_{N\alpha}\}$, i.e. the α-th quantiles of the individual bid distributions. Calculate

$\hat{Y}_i(b_{i\alpha}) = \sum_{j\neq i} \frac{\hat{g}_j(\hat{b}_{i\alpha})}{\hat{G}_j(\hat{b}_{i\alpha})}$ and form the system of equations implied by equation (12):

$$\hat{b}_{i\alpha} - \hat{b}_{j\alpha} = -\frac{\theta_i}{\hat{Y}_i(\hat{b}_i)} + \frac{\theta_j}{\hat{Y}_j(\hat{b}_j)} + \varepsilon_{ij}. \tag{13}$$

Then estimate $\{\theta_i\}_{i=1:N}$ using least-squares.

3.Using the first order condition (12), estimate $F(v)$ as in the risk neutral case.

4 A Simple Adaptive Model

Our previous models assume that bidders "know" the distribution of bids that they are going to face. However, it is entirely possible that the bidders "learn," rather than "know" $Q(b)$, the probability that a bid of b will win the auction. Let h_{it} denote the history of bids observed by the bidder i who submits the bid $b_{(t)}$. Just as the econometrician has to estimate $G(b)$ using the empirical distribution of bids, we assume in this model that bidders estimate $G(b)$ using previously submitted bids. We denote this estimate as $\widehat{G}(b|h_t)$. Given this estimate, bidders choose their bids in order to maximize expected profit $\pi_i(b_i; v_i, \widehat{G}(b|h_t))$ which is equal to

$$\pi_i(b_i; v_i, \widehat{G}(b|h_t)) = (v_i - b_i)\widehat{G}(b|h_t)^{N-1}. \tag{14}$$

In the experiment we consider, bidders were told after each auction what their opponents' bids were, so that

h_t is known to these bidders.

The first order condition for maximization in the learning model is then

$$\hat{v}_{it} = b_{it} + \frac{\hat{G}(b_{it}|h_{it})}{\hat{g}(b_{it}|h_{it})(N-1)} \,. \tag{15}$$

where $\hat{v}_{it}$ is the valuation that rationalizes the t^{th} bid. As in the risk neutral Bayesian-Nash equilibrium model, one can estimate $\hat{G}(b_{it}|h_{it})$and $\hat{g}(b_{it}|h_{it})$ nonparametrically and then estimate $F(v)$ using kernel methods.

The model is above is admittedly a very simple stab at formalizing the intuition that in many real life situations agents learn to play the game correctly through experience. However, if agents form beliefs like econometricians, as in the modeling approach of Sargent (1993), it will not be a bad approximation to many experimental settings, or auction markets where there is a lot of repeated interaction by the same set of actors, and where previous bids are publicly observeable.[6],[7] However, it is worth noting that the data requirements to estimate such a specification model will be heavy, since the economist needs to recreate the information available to the agents at the time of their bidding decision, rather than simply postulating that the agent has rational expectations about the bid distribution she is about to face. Furthermore, there is very little empirical evidence that having better data on past bids allows for better, or different bidding decisions.[8]

[6] Such as many procurement auctions, electricity and Treasury auction markets. In deregulated electricity markets, for example, the same set of producers bid to supply electricity every day, for every hour of the day. Of course, repeated interaction in these settings brings on a whole host of additional concerns such as collusion and other dynamic strategies, which we do not take into account here. This does not mean collusion is not an important issue. In fact, the possibility of collusion considerably restricts the information available to bidders in many applied settings: in many procurement settings alluded to above that only the winning bid, or rough summary statistics of the submitted bids are announced publicly.

[7] It is important to note that the spirit of the analysis in Sargent (1993) and the majority of the learning literature differs from this exercise. Here, we form an econometric estimator using these theories. The literature tend to be concerned with theoretical properties of models where agents look backwards to form beliefs, such as whether standard equilibrium concepts can be supported as limiting behavior in a model with learning.

[8] Most empirical studies on this issue focus on whether more experienced bidders make better, or at least,

5 The Logit Equilibrium Model

Some recent research in experimental economics has made use McKelvey and Palfrey's (1995) QRE model to reconcile deviations from Nash equilibrium predictions. A commonly used variant of QRE is the logit equilibrium model. Analogous to widely used discrete choice models, in the logit equilibrium, a bidder's payoff is the sum of her risk neutral vNM utility, (2), and an idiosyncratic shock which is i.i.d. extreme value. An equilibrium in this model is a distribution of bids that is consistent with maximization for each agent.

In the logit equilibrium model, the set of possible valuations, v_i, and bids, b_i, is assumed to be large, but finite. Let $\mathcal{B} = \{b_1, b_2, ..., b_{\#B}\}$ represent the set of bids agents can choose and let $\mathcal{V} = \{v_1, ..., v_{\#V}\}$ be the set of possible valuations. Unlike the Bayes-Nash model, agents will not use pure strategies in the logit equilibrium. A (symmetric) strategy $\mathbf{B}(b|v)$ is a measure that assigns a probability to every bid b conditional upon a valuation. In order for the strategy to be a well-defined probability measure

$$\text{For all } v \in \mathcal{V} \text{ and } b \in \mathcal{B},\ \mathbf{B}(b|v) \geq 0. \tag{16}$$

$$\text{For all } v \in \mathcal{V}, \sum_{b\in\mathcal{B}} \mathbf{B}(b|v) = 1. \tag{17}$$

That is, no bid can receive less than 0 probability and, conditional upon any valuation v, the probabilities of all the bids must sum to one.

different bidding decisions. One positive finding by Garvin and Kagel (1991) is that more experienced bidders in common value experiments suffer less from the "winner's curse". In field settings, it is typically very difficult to assess "good" bidding decisions from bad. However, Ockenfels and Roth (2002) and Bajari and Hortacsu (2002) have noted that in eBay auctions, where measures of experience are available, more experienced bidders tend to bid later in the auction, which is closer to what equilibrium models of behavior (in an affiliated value setting) might suggest. However, Bajari and Hortacsu (2002), also report that the level of bids submitted by bidders with varying levels of experience do not appear to differ very much holding characteristics of the auction fixed. Hence, it is not clear whether more "experienced" bidders on eBay enjoy higher profits.

If all agents follow the bidding strategy $\mathbf{B}(b|v)$, the probability $Q(b)$ that player i wins the auction with a bid of b satisfies:[9]

$$Q(b) = \left[\sum_{v} \sum_{b'<b} \mathbf{B}(b'|v) f(v) \right]^{N-1} . \tag{18}$$

The term inside the bracket is the probability that a player submits a bid less than b. Since bids are independent, the probability of winning with a bid of b is the term inside the bracket raised to the power $N-1$.

In a Bayes-Nash equilibrium, the utility to bidder i from bidding b_i with a value of v_i is $\pi(b_i; v_i) = (v_i - b_i) * Q(b_i)$. In this logit equilibrium model, let $\widehat{\pi}(b_i; v_i)$ be the utility that the agent i receives from bidding b when she has a valuation v_i; this is a sum $\pi(b_i; v_i)$ and $\varepsilon(b_i, v_i)$:

$$\widehat{\pi}(b; v_i) \equiv (v_i - b_i) * Q(b_i) + \varepsilon(b_i, v_i) = \pi(b_i; v_i) + \varepsilon(b_i, v_i) . \tag{19}$$

The logit equilibrium model generalizes the Bayes-Nash model by including the term $\varepsilon(b_i, v_i)$ in an agent's payoffs. One can interpret $\varepsilon(b_i, v_i)$ as the agent's optimization error. The logit equilibrium model generalizes the Bayes-Nash model by allowing for the possibility that agents fail to perfectly optimize due to the influence of $\varepsilon(b_i, v_i)$ on decision making.

The decision process for a single agent can be thought of as follows: first, each bidder i learns her private information v_i. Second, for every $b_i \in \mathcal{B}$, bidder i draws an error term $\varepsilon(b_i, v_i)$. Finally, each bidder chooses the bid $b \in \mathcal{B}$ which maximizes $\widehat{\pi}(b_i, v_i)$, her expected profit plus $\varepsilon(b_i, v_i)$.

Assume that $\varepsilon(b_i, v_i)$ has a cumulative distribution function $F(\epsilon) = \exp(-\exp(-\lambda\epsilon))$. This distribution

[9] We assume that a bid strictly less than all other bidders is required to win the auction, and thus avoid consideration of ties between bidders. This is to simplify exposition of the problem and the computations. With a sufficiently large set of types and bids, this assumption will not change our results.

has a mean of $\frac{\gamma}{\lambda}$, where γ is Euler's constant (0.577), and variance $\frac{\pi^2}{6\lambda^2}$. Note that λ is proportional to the precision (the inverse of the variance).

Let $\sigma(b_i; v_i, \mathbf{B})$ be the probability that agent i bids b_i conditional on a value draw v_i and that the $N-1$ other agents bid using the strategy $\mathbf{B}$. By well known properties of the extreme value distribution, it follows immediately that:

$$\sigma(b_i; v_i, \mathbf{B}) = \frac{\exp(\lambda \pi(b_i; v_i, \mathbf{B}))}{\sum_{b' \in \mathcal{B}} \exp(\lambda \pi(b'; v_i, \mathbf{B}))}. \tag{20}$$

An equilibrium is a bidding function $\mathbf{B}(b|v)$ that is a fixed point of (20), that is $\mathbf{B}(b|v_i) = \sigma(b_i; v_i, \mathbf{B})$.

As $\lambda \rightarrow 0$, the variance of the error term becomes infinite so that $\pi(b_i; v_i)$ is swamped by the error term $\varepsilon(b_i, v_i)$ in $\widehat{\pi}(b; v_i)$. If $\lambda \rightarrow \infty$, then the variance of the error term tends toward zero and the equilibrium of the game converges to a Bayes-Nash equilibrium. Therefore, the logit equilibrium nests Bayes-Nash equilibrium as a special case when $\lambda \rightarrow \infty$. The existence of the logit equilibrium is obtained by using fixed point methods such as in McKelvey and Palfrey (1995).

5.1 Structural Estimation

Unlike the case of risk neutral and risk averse Bayes-Nash models, there are no existing techniques for nonparametric estimation of $F(v)$ in the logit equilibrium model. Therefore, we suggest a straightforward parametric approach.

Let $F(v|\theta)$ denote the distribution of private information conditional on a vector of parameters θ. Let $\widehat{Q}(b)$ be an estimate of $Q(b)$, the probability of winning the auction with bid b. Let $p(b|\theta)$ denote the probability of the bid b given θ. Given $\widehat{Q}(b)$,

$$p(b|\theta,\lambda;\widehat{Q}) = \int_{\underline{v}}^{\bar{v}} \sigma(b_i|v)f(v|\theta)dv = \int_{\underline{v}}^{\bar{v}} \frac{\exp(\lambda(v-b_i)\widehat{Q}(b_i))}{\sum_{b'\in\mathcal{B}}\exp(\lambda(v-b')\widehat{Q}(b'))}f(v|\theta)dv. \tag{21}$$

Our approach for estimating the logit equilibrium model can be summarized as follows:

1. Given a data set of T bids, $b_{(1)}, ..., b_{(T)}$, form an estimate $\widehat{Q}(b)$ of $Q(b)$.
2. Estimate θ and λ using maximum likelihood, where the likelihood function for θ and λ is

$$L(\theta,\lambda) = \prod_{t=1}^{T} p(b_t|\theta,\lambda;\widehat{Q}).$$

Observe that our two-step procedure eliminates the need to compute the equilibrium, as in McKelvy and Palfrey (1995). This greatly reduces the computational complexity of estimating the model.

It is also worth noting that the logit equilibrium model will not be identified with bid data from the repetition of a single auction. As can be seen in equation (3), the valuations are just identified from the distribution of bids in the risk neutral Bayes-Nash equilibrium model. This suggests that we can never reject the hypothesis that $\lambda = \infty$ (so that $\varepsilon(b_i, v_i)$ is everywhere zero) and that the valuations satisfy (3). Obviously, the logit equilibrium model is not identified in this case and some additional source of variation will be required to identify both λ and $F(v)$. Fortunately, in our experimental data, the number of bidders varies from 3 to 6, but the distribution of private information remains fixed. If we incorporate this information into the likelihood function, then identification of both λ and $F(v)$ is possible. We recognize that the econometrician may not have the luxury of making this assumption in the field. However, comparing the logit equilibrium to alternative models is not possible without this assumption.

Goeree, Holt, and Palfrey (2002) find that a QRE model with risk aversion seems to fit laboratory experiments in first-price sealed-bid auctions well. For instance, if we use the CRRA specification of the previous section, equation (21) becomes

$$p(b|\theta,\lambda;\widehat{Q}) = \int_{\underline{v}}^{\bar{v}} \sigma(b_i|v) f(v|\theta) dv = \int_{\underline{v}}^{\bar{v}} \frac{\exp(\lambda(v-b_i)^{\alpha_i}\widehat{Q}(b_i))}{\sum_{b'\in\mathcal{B}}\exp(\lambda(v-b')^{\alpha_i}\widehat{Q}(b'))} f(v|\theta) dv. \tag{22}$$

We attempted to estimate this specification using our data, but found that the fit was extremely poor. In Goeree, Holt, and Palfrey (2002), the authors treated the valuations v as data and estimated λ and α_i. In our exercise, we also need to estimate θ, the parameters characterizing the distribution of private information. Our task is much more demanding, and therefore, it should not be surprising that (22) does not work particularly well in our application, although it appears to work well in the application considered by Goeree, Holt and Palfrey (2002).

6 Results

The data set we use was provided to us by John Kagel, and contains a series of IPV first-price auction experiments conducted by Dyer, Kagel, and Levin (1989). Their paper contains a detailed explanation of the experimental setup, but here we repeat some of the vital information. There were 3 experimental runs with 6 different subjects participating in each. The subjects were recruited from MBA students at the University of Houston. In these experiments, bidders were assigned i.i.d. valuations drawn from a uniform distribution on $[\$0,\$30]$, and, in the event they won the auction, they were paid their assigned valuation minus their bid. Each subject participated in 28 auctions, during the course of a two hour experimental run. Data from the first 4 runs of the experiments are excluded. This leaves us with 3 runs of 24 auctions.

A novel feature of this experiment was that the bidders were faced with two possibilities as to how many competitors they were going to face: with probability $1/2$, the market they competed in contained $N=3$ bidders and with probability $1/2$, $N=6$ bidders. Bidders were asked to submit two "contingent" bids and one "non-contingent" bid. After the bidders submitted their three bids, a coin was tossed, first, to determine whether the "contingent" or "non-contingent" bids would be used in determining the winner. A second coin toss determined whether $N = 3$ or $N = 6$. If the "contingent" treatment was selected, the first "$N = 3$ contingent" bid was used if $N = 3$, and the "$N = 6$ contingent" bid was used if $N = 6$. That is, if bidder A, with valuation $25, bid $20 contingent on there being 3 bidders, and $18 on the contingency that there would be 6 bidders, and if it turned out that the coin flip favored $N = 3$, bidder A would pay $20 upon winning. After each auction, bids and corresponding private values were posted on a blackboard for the bidders to see.

Throughout the rest of this paper, we will ignore the "non-contingent" bids and focus on the "contingent" bids.[10] This allows us to compare strategic decisions made for the $N = 3$ contingency against the $N = 6$ contingency, keeping the decision-maker and her private valuation constant. We acknowledge that this is a non-standard experimental setup, since bidders are supposed to make three simultaneous bidding decisions instead of one. This may change their response and perhaps increase the frequency of mistakes. However, as we will briefly argue below, as was argued in Dyer, Kagel, and Levin (1989), the main behavioral patterns observed here replicate those seen in other first-price auction experiments. The main advantage of this setup is that there is variation in the number of bidders, holding the distribution of private information fixed. As discussed earlier, this is necessary to identify some of our models. Also, this will give us insight into some of the strengths and weaknesses of these different estimators.

6.1 Full Information Inference

Before trying to recover bidders' valuations using the structural estimation techniques discussed above,

[10] We should note that Dyer, Kagel and Levin (1989) modeled the "non-contingent" bids as being submitted in an environment where there is uncertainty in the number of competitors, and tested the comparative static implications of this uncertainty. Their comparison of revenues across treatments with certain and uncertain number of competitors was consistent with the presence of risk aversion.

we discuss how some of the models described above are able to fit the observed experimental behavior, using information on bidders' actual valuations. As first shown by Vickrey (1962), the symmetric risk neutral Nash equilibrium strategy of the bidders in this symmetric IPV first-price auction with uniformly distributed valuations is:

$$b(v) = \frac{N-1}{N} v. \tag{23}$$

When observed bids were compared with the predicted bids ($\frac{2}{3}v$ in the 3 bidder case, $\frac{5}{6}v$ in the 6 bidder case), Dyer, Kagel, and Levin (1989) found that the observed bids were indeed above the predicted bids.[11] Figures 1(a) and 1(b) plot observed vs. predicted bids. The mean (absolute) deviation of observed bids from risk neutral Bayesian-Nash equilibrium bids when $N = 3$ was \$2.55 with a standard deviation of \$ 1.64. In the 6 player game, the mean deviation is \$0.94 with a standard deviation of \$0.78.

The observed overbidding phenomenon suggests that accounting for risk aversion may help improve the fit of the Bayesian-Nash equilibrium model, both from the experimental standpoint where bidder valuations are observed, and from the structural estimation standpoint. If bidders possess CRRA utility functions, Cox, Walker and Smith (1988) show that the Bayesian-Nash equilibrium bid functions with uniformly distributed private values are given by:

$$b_i(v) = \frac{N-1}{N-r_i} v.$$

where $r_i = 1 - \theta_i$ is the risk aversion coefficient of bidder i. Given this functional form, we can estimate r_i by running the regression:

[11] Similar overbidding has been reported in many other studies, including, for example, Cox, Smith, and Walker (1988) and Harrison (1989). See Kagel and Roth (1997), Chapter 7 for a detailed survey.

$$b_{it} = \alpha_i + \beta_i v + \varepsilon_{it}.$$

and calculate $r_i = N - \frac{N-1}{\beta_i}$.

Table 1 reports the results of this regression, run separately for all bidders in the data set, treating the $N = 3$ and $N = 6$ bids separately. The confidence intervals for θ_i are calculated using the delta method. First, observe that the estimates of α_i are in almost all cases statistically indistinguishable from zero, passing a weak test of the theory. Also observe that the risk aversion estimates for the $N = 3$ contingent bids are near previously reported estimates of risk aversion from other auction experiments, which are around 0.5 (see, for example, Cox, Walker, Smith (1988), Goeree, Holt and Palfrey (2002)).

However, the risk aversion estimates using $N = 6$ contingent bids are quite different, and, in many cases, statistically significantly so. In several cases, the estimated risk aversion for the $N = 6$ are not statistically different from 0, and in general, most estimated coefficients are lower. This result is suggested by figure 1(b), where we observe much less overbidding in the $N = 6$ bids. This discrepancy between the two sets of risk aversion estimates is inconsistent with the CRRA specification, hence we take this as suggestive evidence that the CRRA asymmetric risk aversion specification is not necessarily the *a priori* preferable econometric specification.[12]

Another possible explanation for the observed deviations from the risk neutral Bayes-Nash equilibrium might be that the costs of deviating from equilibrium behavior in the game are small, causing the bidders to optimize imperfectly or even randomly within a flat region of the best response function. This explanation

[12] The discrepancy could potentially be explained away with a more general specification for risk preference, in which relative risk aversion also depends on the level of the payoffs (hence yielding lower implied relative risk aversion with the smaller payoffs in the $N = 6$ case). However, we are not aware of any prior empirical or experimental work using this more general specification.

was first suggested by Harrison (1989).

Indeed, we find that the average difference, per auction, between the expected payoff of the bidder from his observed action and his Nash equilibrium action is not very large: \$0.36 (\$0.45) for $N = 3$, and \$0.05 (\$0.10) for $N = 6$. Figure 2 plots the expected foregone earnings of a bidder from deviating from the Nash equilibrium bid for three different valuations, $v = 3, 12$, and 24 for the $N = 3$ bidder case (the corresponding Nash equilibrium bids are $b^* = 2, 8$, and 16 respectively). As can be observed from this figure, the expected cost of deviating from the Nash equilibrium strategy is not very large, even for large deviations. For a bidder with $v = 24$, the cost of over-bidding by \$4 is about \$1. For a bidder with a small valuation, say $v = 3$, the expected monetary loss from submitting any bid that is below the value draw is negligible.

However, a closer look at the bidding data reveals an inconsistency with the Harrison critique: deviations from Nash equilibrium seem to be smaller in the $N = 6$ case. In figure 3, we plot the foregone payoff for deviations from Nash equilibrium for a bidder with $v = 24$. From this figure, it is clearly more costly to bid above the Nash strategy when $N = 6$ than when $N = 3$. On the other hand, it is more costly to bid below the Nash strategy when $N = 3$. If costs of deviation regulate the magnitude of errors in decision-making, this figure suggests that we should observe less overbidding when $N = 6$ as opposed to $N = 3$, but more underbidding. However, in the data, the degree of underbidding is not appreciably less for the $N = 6$ case, which can not be accounted for by a straightforward foregone payoff argument.

Our conclusion from these observations is that, even with full access to the experimentally assigned private valuations, it is difficult to tell which model can best explain the observed behavior, and hence constitutes the *ex ante* best structural econometric specification. Therefore, we now proceed to rank the different specifications on the basis of their ability to recover the structural parameters.

6.2 Symmetric Risk Neutral Bayesian-Nash Case

To estimate the symmetric risk neutral Bayesian-Nash model, the strategy described in section 2.1 is followed. When estimating the density of bids, $g(b)$, a Normal kernel is used, with the bandwidth set equal to Silverman's 1.635 times sample standard deviation rule-of-thumb.[13] The cdf $G(b)$ is estimated using a simple counting estimator. To recover $\widehat{v}_{i,t}$, the valuation associated with the bid $b_{(t)}$, equation (5) is used. All of the bids were pooled across runs of the experiment. Figure 4 plots the estimated bidder valuations against the assigned valuations, where the bids from all 3 runs of 23 auctions are pooled together to estimate $g(b)$ and $G(b)$.

Observe that in the $N = 3$ case, the estimated valuations are noticeably above the valuations assigned to bidders in the experiment. This is consistent with the overbidding phenomenon discussed above. However, when we construct the Kolmogorov-Smirnov test statistic, KS, where

$$KS = \max |\hat{F}^{RN}(v) - F^{ACT}(v)|.$$

Tables 2 reports the K-S statistics and the associated p-values. We find that in the $N = 3$ case we can reject the equality of the estimated and true distribution of valuations, but that we cannot in the $N = 6$ case.[14]

In Table 2, we also compare the actual and estimated average bid shading undertaken by the bidders, as a second metric. In both cases, the valuation estimates given by equation (5) lead to bid shading estimates that are lower than actual. In the $N = 3$ case, the average bid is shaded 30% while the econometric estimate of

[13] We checked robustness of the kernel estimates by redoing the estimation parametrically with a uniform specification.

[14] In these hypothesis tests, we do not account for the fact that the valuations are themselves random. This would not be difficult to do for the symmetric risk neutral, asymmetric risk neutral and risk averse models. However, in the learning model, the valuations are a function of the entire sequence of previous bids. We are unaware of methods that could be applied to compute standard errors for the valuations, let alone conduct hypothesis testing similar to Table 2. As a result, we abstract away from this issue.

this amount is 18%. Observe, however, that in the $N = 6$, the difference between observed and estimated average bid shading percentages are quite close. From these results, it appears that the symmetric risk neutral Bayesian-Nash equilibrium model performs reasonably well in the $N = 6$ case, but not as well in the $N = 3$ case.

6.3 Asymmetric Risk Neutral Case

Next, we apply the asymmetric estimator of section 2.2. To estimate the bidder level bid distributions and densities, $G_j(b)$ and $g_j(b)$, we once again use a Normal kernel with Silverman's rule-of-thumb bandwidth. In figure 5, we plot the histogram of the estimated valuations from both the symmetric and the asymmetric model. In the $N = 3$ case, it is clear that the asymmetric estimator does a considerably better job of estimating the distribution of true valuations (which is uniform on the interval $[0, 30]$) than the symmetric model. The symmetric model does a poor job of matching the true distribution at the right end of the distribution. As we can see in figure 5, the symmetric model estimates a significant mass of valuations between 30 and 40. In the $N = 6$ case, both estimators seem to do a reasonably good job of recovering the distribution of valuations. Not surprisingly, in Table 3, we find that the value of the K-S statistic has decreased compared to the symmetric risk neutral benchmark in the $N = 3$ case. In both the $N = 3$ and $N = 6$ case, we cannot reject the hypothesis that the true distribution and the estimated distribution are identical.

These results demonstrate that allowing bidders to have different valuations can make a substantial improvement in the estimates of the structural parameters when the number of bidders is small. The asymmetric model implies considerably more sophistication on the part of bidders than the symmetric model. In the asymmetric model, the bidders must take into account when optimizing that each participant uses a different bidding strategy. One interpretation of our results is that the symmetric model produces poorer estimates because it is an overly simplistic representation of bidder's payoffs. Experimental subjects act as if they understand that other subjects do not behave identically.

In Table 4(a) and 4(b), to explore the extent of heterogeneity in how bids are submitted, we run the following regression:

$$b_{i,t} = constant + \alpha v_{i,t} + \beta_i \cdot d(i) \cdot v_{i,t} + \varepsilon_{i,t}. \tag{24}$$

In equation (24), the bid, $b_{i,t}$, is modeled as a function of the valuation assigned in the experiment $v_{i,t}$, and a bidder i specific dummy variable $d(i)$ interacted with i's valuation. In Table 4(a), when there are $N = 3$ bidders, 6 of the 17 bidder specific dummy variables are significant at conventional level. While many of the coefficients are statistically different, their magnitudes are not particularly large, with the notable exception of bidder 12. In Table 4(b), we see that there is less asymmetry in the $N = 6$ bidder case. Only 4 of the 17 interaction terms are significant at conventional levels. Also, the magnitudes of the bidder specific coefficients are smaller than the $N = 3$ case. For both the 3 and 6 bidders case, we reject the hypothesis that the β_i are equal, i.e. that players use the same bid functions.

In Tables 5(a) and 5(b), we explore whether the estimated valuations show significant differences across bidders by running the following regression:

$$\widehat{v}_{i,t} = constant + \alpha v_i + \beta_i \cdot d(i) \cdot v_{i,t} + \varepsilon_{i,t}. \tag{25}$$

In equation (25), we regress our estimate of i's valuation, $\widehat{v}_{i,t}$ on the valuation assigned in the experiment plus a set of bidder specific dummies interacted with the valuations. When there are $N = 3$ bidders, 8 of the bidder specific interaction are significant at the 5 percent level and 5 are significant when $N = 6$. Once again, many of the magnitudes are not particularly large, except for bidder 12, who acts as if he receives a considerable “thrill” from winning. It is interesting to note that in both the $N = 3$ and $N = 6$ case, bidder 12 valuation is estimated to be 1.23 times the true valuation. Bidder 3 and bidder 17, who have the next largest difference between their estimated and true valuation.

In both the $N = 3$ and $N = 6$ case, we reject the hypothesis that the bidder specific interactions are all equal. Once again, the asymmetries are more pronounced in the 3 bidder case than in the 6 bidder case. It is worth keeping in mind, however, that despite these asymmetries, figure 5 demonstrates that the asymmetric estimator seems to do quite a reasonable job of recovering the distribution of valuations. Also, the regressions in Table 5 demonstrate that the difference between the estimated and true valuations is on average fairly small in economic terms, usually considerably below one dollar.

As an aside, note that by allowing for asymmetries, we can rationalize the overbidding phenomenon without resorting to an explanation based on risk aversion. We have no deep explanation of why these asymmetries should exist. As we discussed in section 2.2, one possible source for these asymmetries is that some participants enjoy the experiments more than others and experience a thrill from winning and hence bid in a very aggressive fashion. Another possibility is that certain bidders are inattentive or less skilled in playing bidding games. The current experimental design does not allow us to sort out these competing theories in a compelling manner. However, the regressions in Tables 4 indicate that a small subset of the players, most notably player 12, bid in a much more aggresive fashion than a typical subject. In making a best response, a rational agent should take this into account. The asymmetric model allows for this possibility while the symmetric model does not.

6.4 Asymmetric Risk Aversion Case

One suggested explanation for the "over-bidding" phenomenon is bidder risk aversion. Therefore, we implement the structural estimation method of Campo (2001), outlined in section 3, to see if estimating bidder level risk aversion coefficients leads to reasonable estimates. We estimate $G_j(b)$ and $g_j(b)$ as in the asymmetric risk neutral case. We then pick 4 equally spaced quantiles between 0.1 and 0.9, and evaluate the system of equations (13). The bidder specific risk aversions $1 - \theta_i$ are estimated using least-squares.

In Tables 6(a) and 6(b), we summarize our estimates of the risk aversion coefficients for the 18 subjects across our 3 experiments. Observe that, except for experiment 1, the risk aversion coefficients estimated with this method are nowhere near the risk aversion coefficients we estimated using the assigned valuations of the bidders. Moreover, the risk aversion coefficients we found for experiments 2 and 3 are all statistically greater than 1.

Despite numerous attempts at modifying Campo's procedure, we were not able to generate reasonable estimates of the risk aversion coefficients from experiments 2 and 3. In order to impliment the second step of Campo's procedure, given the lack of better alternatives, we impose the assumption that $\theta_i = 0.5$ for all bidders. This is the experimental benchmark figure and our results from experiment 1 data yielded risk aversion coefficients near this figure.

Table 7 reports the Kolmogorov-Smirnov statistics for the estimate of $F(v)$ under the assumption that $\theta_i = 0.5$. In both the $N = 3$ case and the $N = 6$ case we are not able to reject the equality of the structural estimates and the true distribution (although just barely when $N = 3$). Also, observe that in the $N = 3$ case, the asymmetric risk neutral model has a smaller K-S statistic than the risk averse model, while when $N = 6$, the opposite is true.

Our conclusions from this section are two-fold: First, we conclude that the econometric specification of Campo (2001) does a good job of uncovering $F(v)$ if we assume that the econometrian knows $\theta_i = 0.5$. However, in the field, the economist may not have access to exogenous estimates of risk aversion. Second, we have found Campo's procedure not to be very robust in this experimental setting. It is quite puzzling why this procedure yielded very high risk aversion coefficients for certain subsets of data. Also, we have found the estimates for these experiments not to be robust to inclusion and exclusion of data points. It is worth noting that the stakes in this experiment were not particularly high. In real world markets where the potential payoffs are much larger, the risk aversion model might work much better.

6.5 Quantal Response Equilibrium

Next, the valuation distribution is estimated assuming bidders are playing the (logit version of the) quantal response equilibrium model. As we discussed in section 4, since the risk neutral Bayes-Nash model is just identified, the logit equilibrium model will typically be under-identified. Therefore, in forming the likelihood function, we assume that the econometrician knows that the distribution of valuations is the same in the $N = 3$ and $N = 6$ cases. Furthermore, we assume that the valuations are distributed $[vlower, vupper]$ so that $vlower$, $vupper$, and λ are estimated using maximum likelihood. We recognize that in many real applications, most econometricians do not have this type of a priori information about the structural parameters.

In Table 8, the parameter estimates are reported for both the $N = 3$ and $N = 6$ cases. Observe that for the $N = 3$ case, the logit equilibrium model grossly overstates the upper boundary of the valuation distribution, and for $N = 6$, the lower boundary of the support is overstated.

Why does this occur? By equation (20), the logit equilibrium model predicts equal probabilities of placing different bids when the payoff from placing these bids is about equal. We computed that with 6 bidders in the auction, the probability that the bidder wins the auction with a bid less than 10 is less than 1/729. Thus, the bidder's expected payoff from any bid between 0 and 10 is very near zero. Not surprisingly, the logit equilibrium predicts that a bidder with valuation 1 is indifferent between placing any bid between 0 and 10. Hence, a wide range of bids can be "rationalized" by a given valuation, and vice versa. Therefore, the relationship between bids and valuations is not particularly tight.

We found two other unappealing features of the logit equilibrium estimates. First, although the logit equilibrium model assigns positive probability to every action, we do not see bidders playing dominated strategies (nobody bids above their value). This violates the *full support* assumption regarding the distribution of the error term. Second, the estimated magnitude of λ, the variance of the logit error term, changes with the number of bidders in the game. In the basic logit equilibrium model, there is no accounting for how different types of games lead to different types of errors. Even if we let λ be a (monotonic) function of the expected payoff (that is, the bidder tries harder to resolve his payoff uncertainty if there is more at stake), this would predict that decision errors should be larger in the $N = 6$ game. This is not the case. However, we find that instead of bidding anything, bidders essentially bid their valuation given that they know they have a very low probability of winning. Therefore, the predictions of the equilibrium model actually become more consistent with the theory as the number of bidders increases. This is the opposite of what the logit equilibrium model suggests.

6.6 Learning Model

To see if the imperfect overlap between the bidders' information set and the econometrician's information set is the source of any estimation bias, the distribution of bidder valuations is estimated using the learning framework described in section 5. Once again, following the estimation exercise for the symmetric risk neutral case, Normal kernel estimators for the density and distribution of bids are used. However, this time when estimating $\hat{G}(b_{it}|h_{it})$, the estimation sample was restricted to the bids previously observed by the bidders. Recall that the experimenters wrote down all the realized bids on a blackboard for the bidders to see.

Figures 6 and 7 plot the estimated valuations against the observed valuations. Unlike the symmetric risk neutral Bayes-Nash model, the estimated valuations are lower than the true valuations in the $N = 3$ case and higher than the true valuations in the $N = 6$ case. Unlike the rational models or the QRE models, where data from many auctions is pooled to estimate $\hat{G}(b)$, in the learning model, only the past history of observations in a given run enter $\hat{G}(b|h_{it})$. As a result, the estimates of $\hat{G}(b_{it}|h_{it})$ may differ substantially from run to run. The learning model predicts that this should have a strong influence on bidding behavior. However, figures 6 and 7 suggest that both low and high valuation bidders appear to submit quite different bids than those predicted by the learning model. By comparison, the deviations from equilibrium in the symmetric risk neutral Bayes-Nash model appear to be more stark when the number of bidders is small and the valuations are high.

7 Summary Results

Tables 9 and 10 summarize our results regarding the "closeness" of the estimated distribution of valuations and the true distribution of valuations. To assess the robustness of the Kolmogorov-Smirnov distance measure between $F(v)$ and $\hat{F}(v)$, we computed two additional distance metrics, L^1 and L^2, defined as:

$$
\begin{aligned}
L^1(F,\hat{F}) &= \int |F(v) - \hat{F}(v)| dv \\
L^2(F,\hat{F}) &= \left(\int (F(v) - \hat{F}(v))^2 dv \right)^{1/2}
\end{aligned}
$$

Also, in figures 8 and 9, we compare the estimated and actual cdf for all the models in the $N = 3$ and $N = 6$ bidder cases. Our results suggest that, when $N = 3$, the risk averse and asymmetric risk neutral model do a considerably better job of uncovering the distribution of valuations than the other estimators. In the $N = 6$ case, the risk averse model has the minimum distance in all metrics, except for the K-S statistic. All of the models, except the QRE, seem to do a reasonable job in uncovering the deep parameters.

The fact that estimates from the logit equilibrium came up last is rather eye-opening. In a recent paper, Goeree, Palfrey, and Plott (2002) offer the QRE model with risk aversion as an explanation for the overbidding puzzle. Unfortunately, as we mentioned in section 4, the QRE model with risk aversion had a very poor fit. The estimates of section 6 demonstrate that some sort of error prone behavior is present in the data. However, these results suggest that bidders' errors do not follow easily identifiable patterns. Furthermore, identifying errors in optimization from bidders' private information using standard data sets will probably not be possible without some fairly strong assumptions.

Our interpretation of these results is that when the number of bidders is sufficiently large, most of the methods, except for the QRE, do a good job in uncovering the deep parameters. When the number of bidders is smaller, the results are more sensitive to the choice of method. The risk aversion and asymmetric models are preferred in this setting in part because they allow for heterogeneity in how agents bid. As a practical matter, if we had to choose on model based on our results, we would probably favor the asymmetric risk neutral model. While both the risk aversion and asymmetric model did a good job in both cases, we had to make a strong auxilary assumption about risk aversion, i.e. we set $\theta_i = .5$, in order to estimate the risk averse specification.

Just because the asymmetric risk neutral model appears to dominate other specifications in laboratory data does not imply that it should always be preferred when analyzing field data. The data requirements for estimating the asymmetric risk neutral Bayes-Nash model are heavier than the risk neutral Bayes-Nash model since bidder specific bid distributions must be estimated. In many empirical applications, it is not always possible to estimate bidder specific bid distributions very precisely when there are few observations per bidder and the bid distributions depends on co-variates.

8 Conclusion

In this paper, we have compared four alternative models of bidding behavior in the first-price sealed-bid auctions: 1) Bayesian-Nash equilibrium with risk neutral bidders, 2) Bayesian-Nash equilibrium with (asymmetric) risk averse bidders, 3) Quantal response equilibrium (with symmetric risk neutral bidders) and 4) "best-response" bidding (by symmetric risk neutral bidders). Despite the fact that many economists express great skepticism about the plausibility of the rationality assumptions used in structural estimation, in our model comparison exercises, the "rational" models give better estimates than the "behavioral" models. While there are clearly deviations from rational behavior, the existing "behavioral" models in our experiments do not generate more compelling estimates of the structural parameters.

Despite the fact that, in our chosen metric, some models perform better than others, it is clear that all of the models have unique limitations. Also, just because one model dominates others in the metrics that we have chosen, it does not follow that this model is the best to apply in all possible applications. If one is willing to accept the somewhat controversial assumption that behavior in the lab is indicative of behavior in the field, then this exercise offers the following lessons. First, structural estimates from rational models of bidding behavior are likely to be more accurate as the number of bidders increases. Second, the estimated valuations are more likely to be close to the true valuations for moderate or low valuations. Third, bidders do appear to systematically deviate from rational behavior. While these deviations are not large in monetary terms, it is important to consider the robustness of one's analysis to this type of behavior. Finally, allowing for bidder asymmetries produces better estimates than a model that assumes all bidders have identical bid functions.

9 References

[1] Bajari, P. and Hortacsu, A. "The Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions." Stanford University Working Paper (2002).

[2] Bajari, P. and Ye, L. "Deciding Between Competition and Collusion", forthcoming, *Review of Economics and Statistics*.

[3] Bray, M. "Learning, Estimation, and the Stability of Rational Expectations." *Journal of Economic Theory*, April, Vol. 26, No.2, pp.318-339.

[4] Campo, S., "Asymmetry and Risk Aversion within the Independent Private Paradigm: The Case of the Construction Procurements", UNC Working Paper.

[5] Campo, S., Perrigne, I. and Vuong, Q. "Semiparametric Estimation of First-Price Auctions with Risk Averse Bidders", USC Working Paper (2002).

[6] Capra, M., Gomez, R., Goeree, J. and Holt, C. "Learning and Noisy Equilibrium Behavior in a Study of Imperfect Price Competition", *International Economic Review*, 43(3), August 2002, 613-636.

[7] Cho, I.K. and Sargent, T.J. "Learning to be Credible," Stanford University Working Paper (1997).

[8] Cox, J. C., Smith, V. L., Walker, J. M. "Tests of a Heterogenous Bidder's Theory of First-Price Auctions," *Economic Letters*, 1983a, 12 (3-4), 207-12.

[9] Cox, J. C., Smith, V. L., Walker, J. M. "A Test That Discriminates Between Two Models of the Dutch-First Auction Non-Isomorphism," *Journal of Economic Behavior and Organization*, June-September 1983b, 4, 205-19.

[10] Cox, J. C., Smith, V. L., Walker, J. M. "Theory and Individual Behavior of First-Price Auctions," *Journal of Risk and Uncertainty*, March 1988, 1, 61-99.

[11] Cox, J. C., Smith, V. L., Walker, J. M. "Theory and Misbehavior of First-Price Auctions: Comment," *American Economic Review,* 82(5), 1392-412 (1992).

[12] Donald, S. and Paarsch, H.J. "Piecewise Pseudo-maximum Likelihood Estimation in Empirical Models of Auctions", *International Economic Review* v34, n1 121-48 (1993).

[13] Donald, S. and Paarsch, H.J. "Identification, Estimation, and Testing in Parametric Empirical Models of Auctions within the Independent Private Values Paradigm", *Econometric Theory* v12, n3: 517-67 (1996).

[14] Dyer, D., Kagel, J. H., Levin, D. "A Comparison of Naive and Experienced Bidders in Common Value Offer Auctions: a Laboratory Analysis," *Economic Journal,* 99(394), 108-15 (1989).

[15] Elyakime, B., Laffont, J.J., Loisel, O., and Vuong, Q. "First-Price Sealed-Bid Auctions with Secret Reservation Prices", *Annales d'Economie et de Statistique* v0, n34 115-41 (1994).

[16] Flambard V. and Perrigne, I. "Asymmetry in Procurement Auctions: Some Evidence from Snow Removal Contracts", unpublished working paper, University of Southern California.

[17] Goeree, J. and Holt, C. "Stochastic Game Theory: For Playing Games, Not Just for Doing Theory," *Proceedings of the National Academy of Sciences*, 96, September 1999, 10564-10567.

[18] Goeree, J., Holt, C. and Palfrey, T., "Quantal Response Equilibrium and Overbidding in Private-Value Auctions," *Journal of Economic Theory*, Vol. 104, No. 1, May 2002, pp.247-272.

[19] Guerre, E., Perrigne, I., Vuong, Q. "Optimal Nonparametric Estimation of First-Price Auctions," *Econometrica*, 68(3), 525-74 (2000).

[20] Harrison, G. W. "Theory and Misbehavior of First-Price Auctions: Reply," *American Economic Review* 82(5), 1426-43 (1992).

[21] Harrison, G. W. "Theory and Misbehavior of First-Price Auctions," *American Economic Review* 79(4), 749-62 (1989).

[22] Hendricks, K., Pinkse, J. and Porter, R. "Empirical Implications of Equilibrium Bidding in First-Price Symmetric, Common Value Auctions", forthcoming, *Review of Economic Studies*.

[23] Kagel, J. H., Roth, A. E. "Theory and Misbehavior in First-Price Auctions: Comment," *American Economic Review*, 82(5), 1379-91 (1992).

[24] Kagel, J. H., Roth, A. E. (eds.) *The Handbook of Experimental Economics*. Princeton University Press, Princeton (1997).

[25] Laffont, J.J., Ossard, H., Vuong, Q. "Econometrics of First-Price Auctions," *Econometrica*, 63(4), 953-80 (1995).

[26] Laffont, J.J. and Vuong, Q. "Structural Analysis of Auction Data", *American Economic Review* v86, n2 (May 1996), 414-20.

[27] Marcet, A., Nicolini, J.P., and Sargent, T.J. "Convergence of Least Squares Learning in Environments with Hidden State Varialbes and Private Information," *Journal of Political Economy*, February, Vol. 84, No. 1, pp.113-126.

[28] Maskin, E. and Riley, J. "Equilibrium in Sealed High Bid Auctions", *Review of Economic Studies* v67, n3 (July 2000a): 413-38.

[29] Maskin, E. and Riley, J. "Asymmetric Auctions", *Review of Economic Studies* v67, n3 (July 2000b): 493-54.

[30] Marcet, A. and Sargent, T.J. "Convergence of Least Squares Learning Mechanisms in Self Referential Linear Stochastic Models," *Journal of Economic Theory*, 48: 337-368.

[31] McKelvey, R. D., Palfrey, T. R. "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10(1), 6-38 (1995).

[32] Merlo, A., Schotter, A. "Theory and Misbehavior of First-Price Auctions: Comment," *American Economic Review*, 82(5), 1413-25 (1992).

[33] Milgrom, P. R., Weber, R. J. "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50(5), 1089-1122 (1982).

[34] Paarsch, J.J. "Deciding Between the Common and Private Value Paradigms in Empirical Models of Auctions", *Journal of Econometrics*, v51, n1-2, 191-215.

[35] Pesendorfer, M. and Jofre-Bonet, M., "Bidding Behavior in a Repeated Procurement Auction", mimeo, London School of Economics.

[36] Sargent, T.J. *Bounded Rationality and Macroeconomics*. Oxford University Press, New York (1993).

Tables and Figures.

Table 1: Full Information Estimates of Risk Aversion.

EXPERIMENT	SUBJECT	RISK AVERSION N=3	ALPHA N=3	RISK AVERSION N=6	ALPHA N=6
1	1	0.19 (0.05,0.33)	1.59 (0.55,2.64)	-0.35 (-0.42,0.29)	1.54 (0.49,2.60)
1	2	0.56 (0.46,0.65)	0.57 (-0.14,1.27)	0.39 (0.39,0.40)	0.40 (-0.15,0.95)
1	3	0.49 (0.37,0.60)	0.85 (-0.13,1.83)	0.06 (-0.03,0.15)	0.86 (-0.17,1.90)
1	4	0.72 (0.62,0.83)	-0.78 (-1.71,0.15)	0.62 (0.55,0.70)	-0.97 (-1.89,-0.05)
1	5	0.47 (0.38,0.56)	0.43 (-0.32,1.18)	0.14 (0.11,0.16)	0.23 (-0.46,0.92)
1	6	0.63 (0.55,0.72)	-0.09 (-0.90,0.72)	0.29 (0.22,0.36)	0.22 (-0.55,0.99)
2	1	0.55 (0.33,0.78)	-0.79 (-2.70,1.12)	0.44 (0.36,0.51)	-0.89 (-1.48,-0.30)
2	2	0.46 (0.29,0.62)	0.11 (-1.12,1.33)	0.16 (0.07,0.25)	0.12 (-0.89,1.12)
2	3	0.57 (0.49,0.65)	-0.38 (-1.04,0.27)	0.17 (0.12,0.22)	0.05 (-0.52,0.64)
2	4	0.24 (0.12,0.36)	0.49 (-0.38,1.36)	0.00 (-0.02,0.03)	0.58 (-0.08,1.23)
2	5	0.51 (0.43,0.58)	0.53 (-0.14,1.21)	0.04 (-0.01,0.09)	0.31 (-0.31,0.93)
2	6	0.77 (0.72,0.82)	0.07 (-0.34,0.48)	0.59 (0.57,0.61)	0.12 (-0.19,0.42)
3	1	0.53 (0.43,0.62)	0.34 (-0.47,1.15)	-0.18 (-0.32,-0.04)	0.35 (-0.47,1.17)
3	2	0.47 (0.36,0.59)	0.42 (-0.68,1.53)	-0.15 (-0.27,-0.04)	0.75 (-0.17,1.67)
3	3	0.51 (0.41,0.61)	-0.07 (-1.03,0.89)	0.03 (-0.03,0.10)	0.41 (-0.21,1.03)
3	4	0.44 (0.37,0.51)	-0.77 (-1.26,-0.28)	0.45 (0.38,0.52)	-0.54 (-1.09,0.01)
3	5	0.75 (0.69,0.81)	-0.62 (-1.22,-0.01)	0.48 (0.41,0.54)	-0.72 (-1.37,-0.08)
3	6	0.62 (0.46,0.79)	-0.62 (-1.92,0.69)	0.76 (0.70,0.81)	-2.23 (-3.46,-0.99)

Table 2: K-S Statistics and Bid Shading for Symmetric Risk Neutral Model.

	N=3	N=6
K-S Statistic	0.1795	0.0513
p-value	0	0.6732
Mean estimated bid-shading	0.1800	0.1757
Mean true bid-shading	0.2961	0.1897

Table 3: K-S Statistics for Asymmetric Risk Neutral Bidding.

	N=3	N=6
K-S Statistic p-value	0.038 0.6124	0.0821 0.1379

Table 4(a): Bid Function Regressions for *N=3*.

VARIABLE NAME	COEFFICIENT	T-STATISTIC
v	.8120	56.34
d2*v	.03861	1.99
d3*v	.02292	1.29
d4*v	.02583	1.41
d5*v	.01229	0.66
d6*v	.02400	1.36
d7*v	-.04263	-2.38
d8*v	.007188	0.37
d9*v	-.007421	-0.41
d10*v	-.07757	-4.08
d11*v	.01362	0.74
d12*v	.08851	4.77
d13*v	.01626	0.85
d14*v	-.0004639	-0.03
d15*v	-.02229	-1.18
d16*v	-.08636	-4.10
d17*v	.04820	2.65
d18*v	-.007967	-0.39
Constant	.07867	0.67
NOBS	408	
R^2	0.9755	

6 bidders interaction terms are significant at conventional levels.
Maximum value for interaction term is 0.88, i.e. one participant bids 90% of value.
Test that bidder interaction terms are equal to each other is rejected at 1 percent significance level.

Table 4(b): Bid Function Regressions for *N=6*.

VARIABLE NAME	COEFFICIENT	T-STATISTIC
v	.8795	74.49
d2*v	.03991	2.51
d3*v	.001981	0.14
d4*v	-.002570	-0.17
d5*v	-.002647	-0.17
d6*v	.002461	0.17
d7*v	-.02439	-1.66
d8*v	-.0007101	-0.04
d9*v	-.01841	-1.24
d10*v	-.01518	-0.97
d11*v	-.02712	-1.81
d12*v	.04726	3.11
d13*v	-.05060	-3.23
d14*v	-.03584	-2.51
d15*v	-.03027	-1.96
d16*v	-.009182	-0.53
d17*v	-.01028	-0.69
d18*v	-.06023	-3.64
Constant	.08258	0.86
NOBS	407	
R^2	.9248	

4 bidder dummies are significant at conventional levels.
Maximum value for interaction term is 0.047, i.e. one participant bids 93% of value.
Test that bidder interaction terms are equal to each other is rejected at 1 percent significance level.

Table 5(a): Valuation Regressions for *N=3*.

VARIABLE NAME	COEFFICIENT	T-STATISTIC
v	.9894	46.09
d2*v	.03210	1.11
d3*v	.05242	1.98
d4*v	.03214	1.17
d5*v	.01270	0.46
d6*v	.08206	3.11
d7*v	.004972	0.19
d8*v	.05858	2.01
d9*v	.04574	1.69
d10*v	-.06570	-2.32
d11*v	.06502	2.38
d12*v	.2419	8.76
d13*v	.03583	1.26
d14*v	-.008037	-0.31
d15*v	-.007193	-0.26
d16*v	-.13359	-4.26
d17*v	.09866	3.65
d18*v	-.03064	-1.02
Constant	-.05494	-0.31
NOBS	407	
R^2	0.9678	

8 bidders dummies are significant at conventional levels.
1 bidder has estval = 1.23*v.
Test that bidder interaction terms are equal to each other is rejected at 1 percent significance level.

Table 5(b): Valuation Regressions for *N=6*.

VARIABLE NAME	COEFFICIENT	T-STATISTIC
v	1.066	61.83
d2*v	.03516	1.51
d3*v	.01935	0.91
d4*v	-.01076	-0.49
d5*v	-.009359	-0.42
d6*v	.04872	2.30
d7*v	.01853	0.87
d8*v	.03669	1.57
d9*v	.01658	0.76
d10*v	.008362	0.37
d11*v	-.007212	-0.33
d12*v	.1421	6.41
d13*v	-.07024	-3.07
d14*v	-.07075	-3.39
d15*v	-.01392	-0.62
d16*v	-.02892	-1.15
d17*v	.002772	0.13
d18*v	-.09607	-3.98
Constant	.08278	0.59
NOBS	408	
R^2	0.9803	

5 bidders dummies are significant at conventional levels.
1 bidder has estval = 1.23*truval.
Test that bidder interaction terms are equal to each other is rejected at 1 percent significance level.

Table 6(a): Estimated Relative Risk Aversion Coefficient, N=3 Bidders.

SUBJECT	EXPERIMENT 1	EXPERIMENT 2	EXPERIMENT 3
1	0.76 (0.59)	1.84 (0.61)	2.30 (0.60)
2	0.40 (0.64)	1.84 (0.61)	3.68 (0.84)
3	1.09 (0.58)	1.51 (0.58)	2.46 (0.62)
4	0.90 (0.57)	0.80 (0.88)	0.80 (0.63)
5	0.65 (0.65)	1.95 (0.70)	3.69 (0.76)
6	1.15 (0.48)	1.69 (0.50)	1.92 (0.86)

Table 6(b): Estimated Risk Aversion Coefficients, N=6 Bidders.

SUBJECT	EXPERIMENT 1	EXPERIMENT 2	EXPERIMENT 3
1	1.27 (0.60)	2.35 (0.67)	2.47 (0.59)
2	0.98 (0.62)	1.38 (0.54)	4.24 (0.83)
3	1.35 (0.61)	2.00 (0.61)	3.07 (0.64)
4	1.15 (0.59)	1.62 (0.82)	1.47 (0.54)
5	1.05 (0.69)	2.02 (0.74)	3.91 (0.73)
6	1.38 (0.51)	1.78 (0.54)	1.98 (0.71)

Table 7: K-S Statistics for Risk Averse Bidding.

	N=3	N=6
K-S Statistic	0.0949	0.0667
p-value	0.056	0.3406

Table 8: QRE Estimates.

	N=3	N=6
vlower	2.31 (0.83)	8.44 (0.12)
vupper	38.12 (0.45)	31.21 (0.01)
λ	6.21 (1.08)	4.96 (0.14)

Table 9: Comparison of Models: N=3.

METRIC/MODEL	SYMMETRIC RISK NEUTRAL BNE	ASYMMETRIC RISK NEUTRAL BNE	RISK AVERSE BNE*	RISK NEUTRAL QRE	ADAPTIVE LEARNING
L2 - Mean Sq. Err.	0.0134	0.0038	0.0025	0.0147	0.0163
L1 - Mean Abs. Err.	0.1072	0.0306	0.0155	0.1175	0.1323
Kolmogorov-Smirnov stat	0.1795	0.0538	0.0949	N/A**	0.25
Reject F(v) = Fest(v)?**	Yes	No	No	N/A	Yes
P-value of K-S	0.0000	0.6124	0.056	N/A	0.0000

*Calculated by setting risk aversion coefficient equal to 0.5 for all bidders.
**K-S test of equality of distribution rejected for p-value below 0.05.
*** Can not calculate K-S statistic for QRE in analogous manner, since model is estimated parametrically.

Table 10: Comparison of Models: N=6.

METRIC/MODEL	SYMMETRIC RISK NEUTRAL BNE	ASYMMETRIC RISK NEUTRAL BNE	RISK AVERSE BNE*	RISK NEUTRAL QRE	ADAPTIVE LEARNING
L2 - Mean Sq. Err.	0.0051	0.0068	0.0024	0.0148	0.0064
L1 - Mean Abs. Err.	0.0385	0.0531	0.0176	0.1045	0.0497
Kolmogorov-Smirnov stat	0.0513	0.0821	0.0667	N/A**	0.0784
Reject F(v) = Fest(v)?**	No	No	No	N/A	No
P-value of K-S	0.6732	0.1379	0.3406	N/A	0.1553

*Calculated by setting risk aversion coefficient equal to 0.5 for all bidders.
**K-S test of equality of distribution rejected for p-value below 0.05.
*** Can not calculate K-S statistic for QRE in analogous manner, since model is estimated parametrically.

Figure 1: Actual Versus Bayes-Nash Bids, $N = 3$ and $N = 6$.

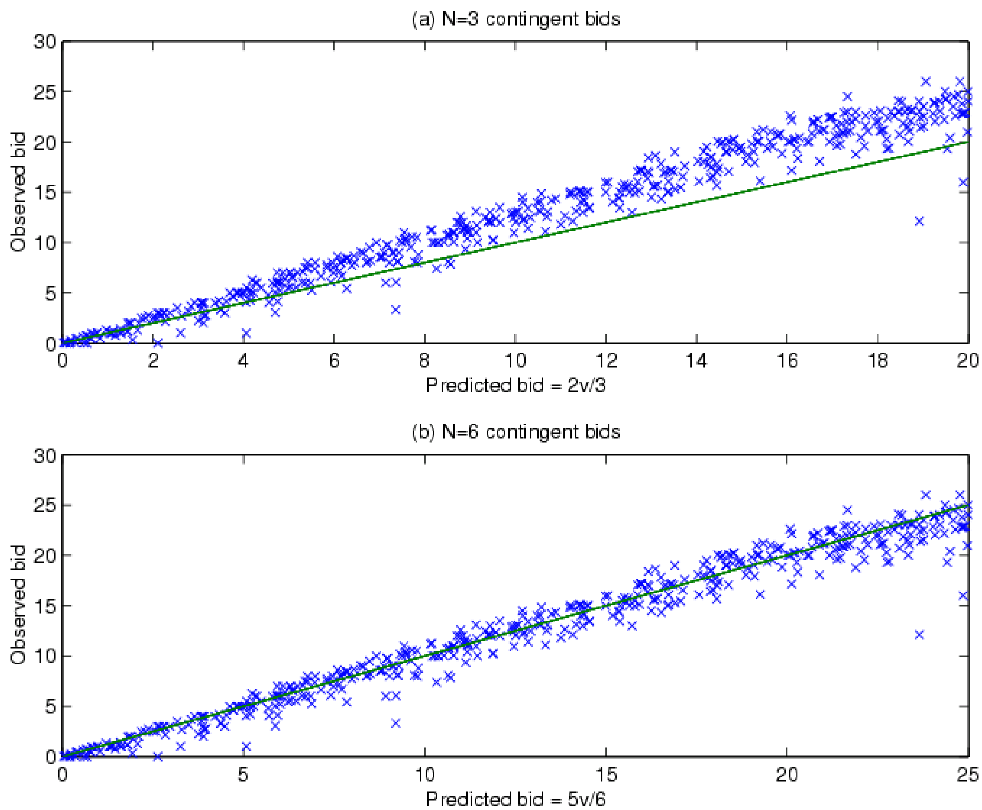

Figure 2: Payoff Forgone by Deviating From Bayes-Nash Bidding Strategy.

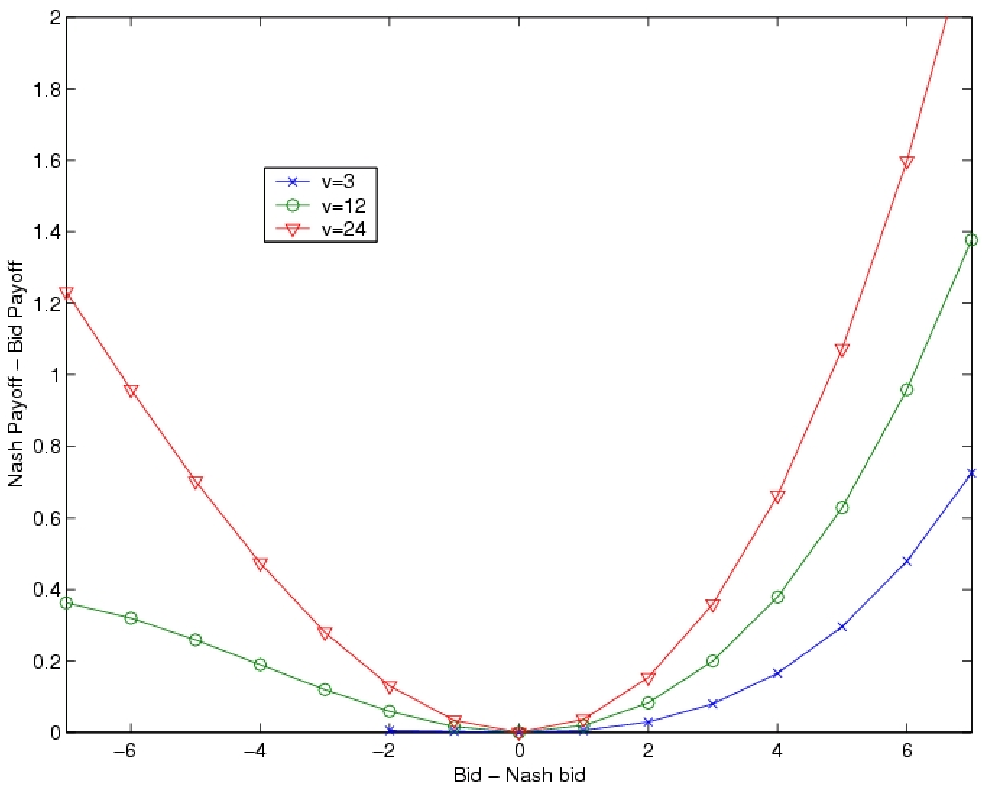

Figure 3: Cost of Deviating From Nash Behavior, 6 Bidder Case Verus 3 Bidder Case.

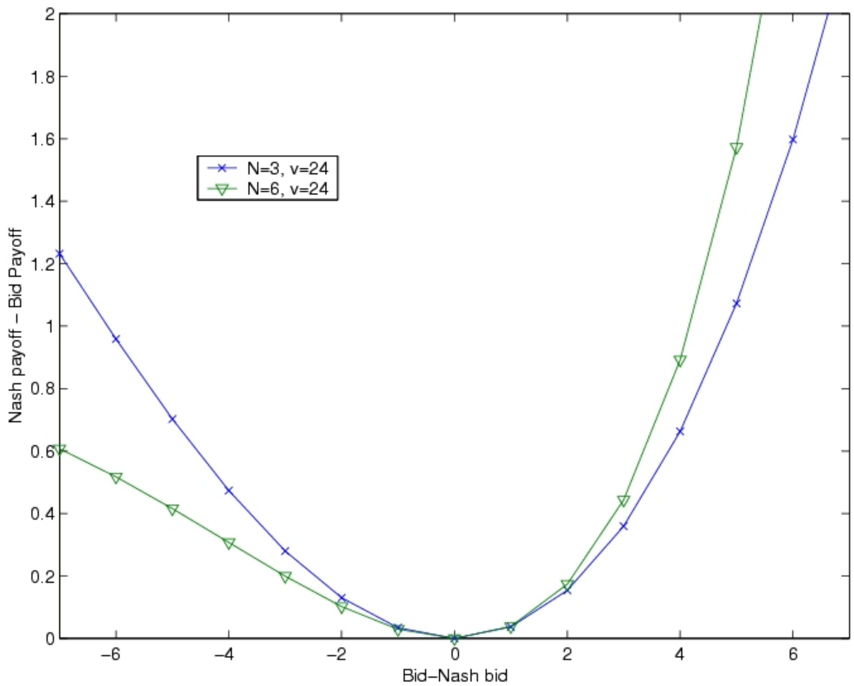

Figure 4: Estimated Versus Actual Valuations for Symmetric Risk Neutral Model, $N = 3$.

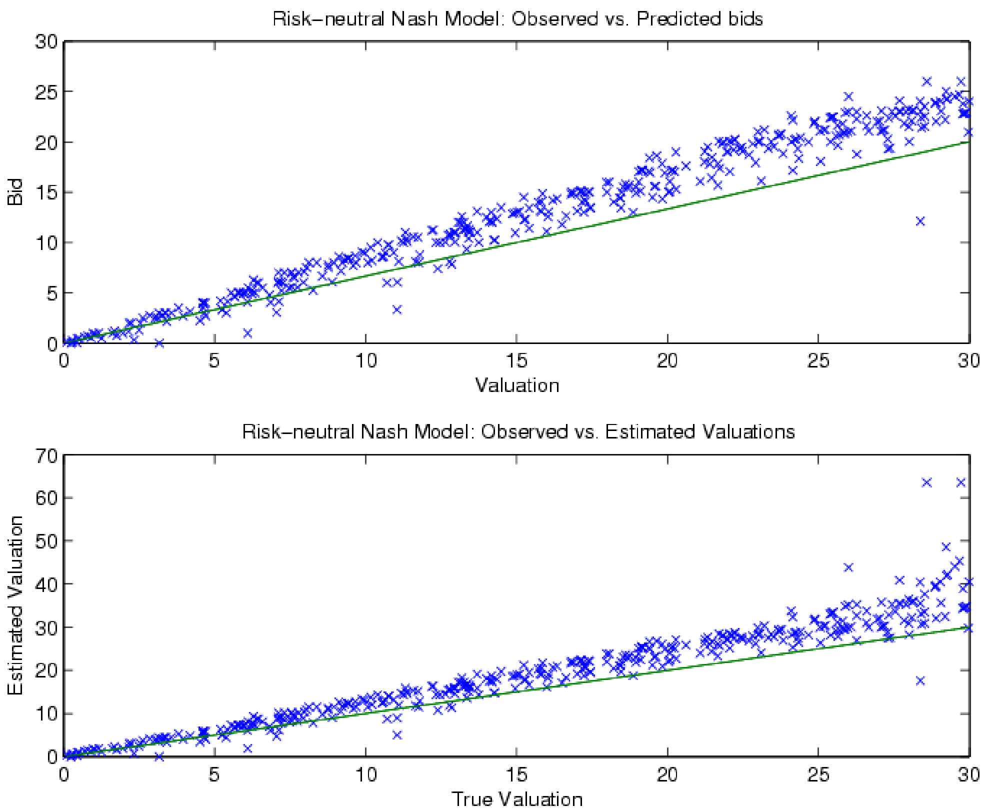

Figure 5: Histogram of Estimated Valuations from Symmetric and Asymmetric Models.

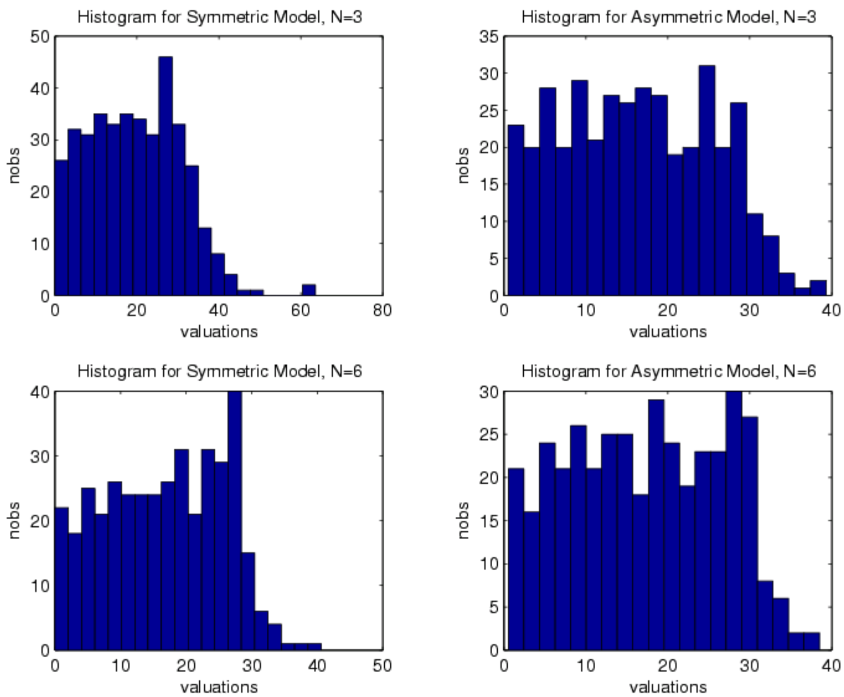

Figure 6: Estimated Versus True Valuations, Learning Model, $N = 3$.

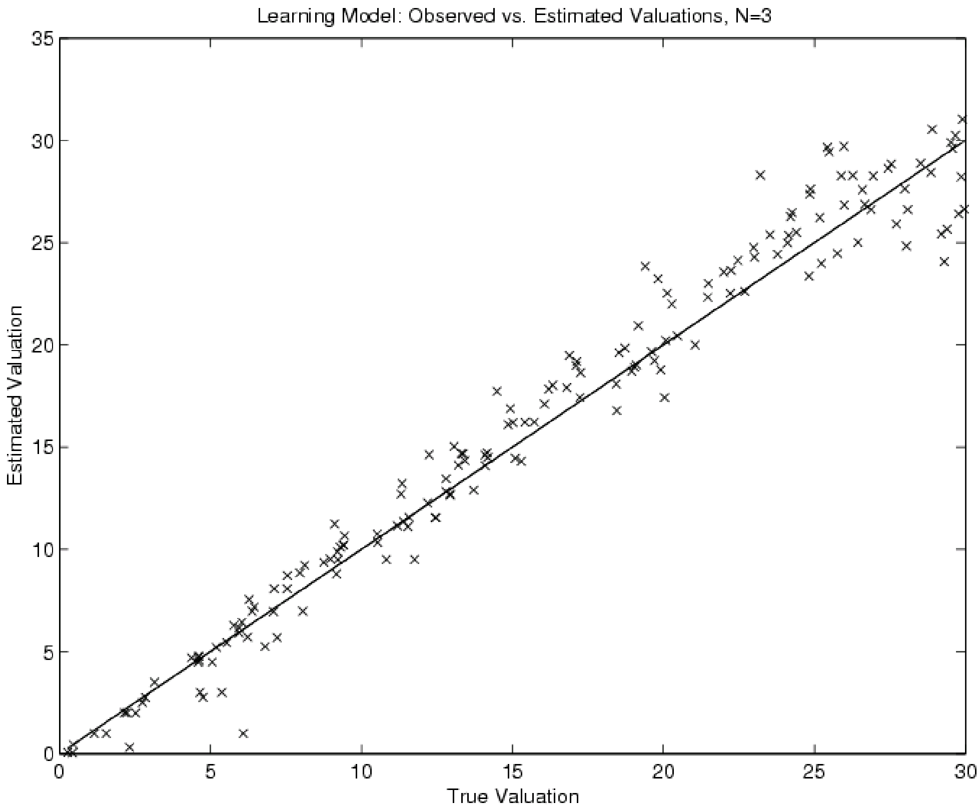

Figure 7: Estimated Versus True Valuations, Learning Model, $N = 6$.

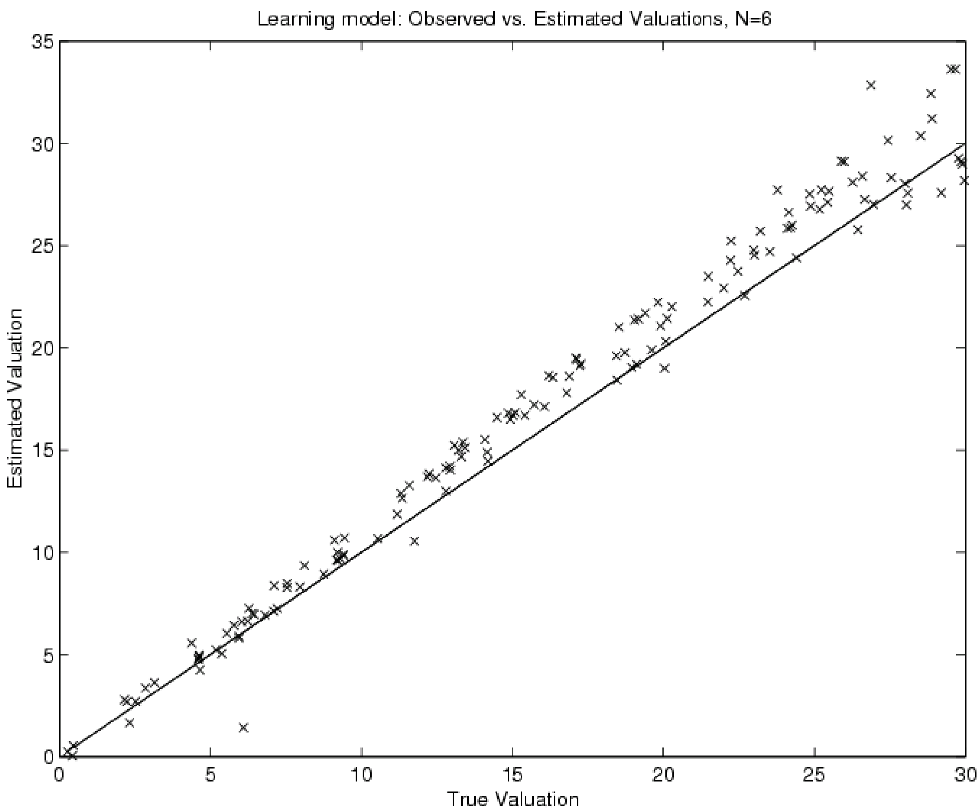

Figure 8: Comparison of True and Estimated CDF's, $N = 3$ Case.

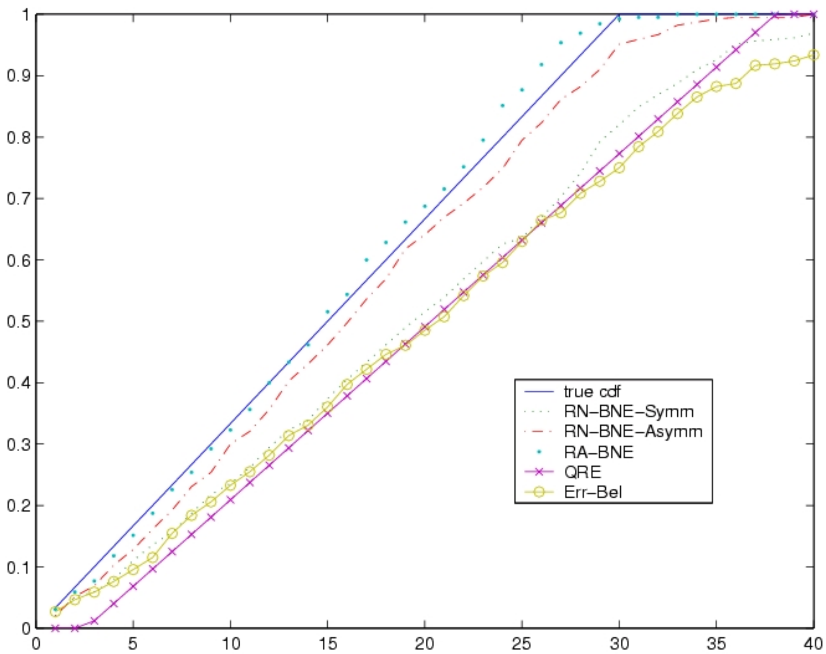

Figure 9: Comparison of True and Estimated CDF's, $N = 6$ Case.