Are Structural Estimates of Auction Models Reasonable? Evidence from Experimental Data.

Patrick Bajari and Ali Hortacsu[1]
University of Michigan and NBER
University of Chicago and NBER

May 21, 2005

Abstract

Recently, economists have developed methods for structural estimation of auction models. Many researchers object to these methods because they find the strict rationality assumptions to be implausible. Using bid data from first-price auction experiments, we estimate four alternative structural models: 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 the risk aversion model is able to generate reasonable estimates of bidder valuations.

---

[1] We would like to thank John Kagel for generously providing us with the data used in this study, and Lanier Benkard, Liran Einav, Han Hong, Joris Pinkse and Peter Reiss for helpful comments. Bajari and Hortacsu would like to acknowledge the National Science Foundation, grants SES-0112106 and SES-0122747 (Bajari) and SES-0242031 (Hortacsu), for generous research support.

1 Introduction

The use of structural econometric models in empirical industrial organization research has become increasingly common. In particular, the empirical analysis of auction data has been transformed by these methods. As pointed out in the influential survey of Laffont and Vuong (1996), auction models appear especially well-suited for structural estimation "...because of the availability of many data sets and the well-defined game forms associated with auctions." In a structural auction model, the economist estimates bidders' reservation values. Given these estimates, one can then study policy counterfactuals, such as changes in reservation price policy and auction format.[2]

Despite the introduction of powerful new methods to estimate these models, many applied researchers are not comfortable with the strict rationality assumptions imposed in the econometric analysis. This skepticism is not without merit. The structural approach is based on three strong assumptions. The first is that bidders' goal is to maximize their expected utility. Second, bidders are able to compute the relationship between their bid and the probability of winning the auction. Third, given their beliefs, bidders are able to correctly maximize expected utility. Even the most ardent supporter of equilibrium behavior might be concerned about imposing such a sophisticated model of behavior in an empirical application.

In the field, it is typically difficult to detect whether bidders are behaving rationally. A number of

---

[2] 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), Guerre, Perrigne and Vuong (2000), Flambard and Perrigne (2001), Campo (2001), Campo, Guerre, Perrigne and Vuong (2000) and Hendricks, Pinkse and Porter (2002).

papers attempt to test *necessary* conditions for rationality. For example, Guerre, Perrigne and Vuong (2000) demonstrate that, in principle, one can test whether the bid function is increasing which is a necessary condition for rationality in private values models.[3] In the field, however, such tests are rendered less powerful by the presence of omitted variables and possible misspecification of the model.

In the field, our ability to test *sufficient* conditions for rationality in first-price auctions is also limited. Some researchers have compared estimates of bidder markups in auctions to actual markups by exploiting accounting data.[4] If the estimated and actual markups were equal, this would be close to a sufficient condition for rationality. While such tests are informative, measures of bidders' private valuations based on accounting data are likely to be imperfect. Accounting costs often do not correspond to economic notions of cost. Moreover, accounting data on markups in auctions is rarely available, making it difficult to test sufficient conditions for rationality in most applications.

In summary, our ability to verify or reject bidder rationality is imperfect in empirical applications. Hence, in most applications, the strong rationality assumptions of structural auction models will be *identifying assumptions*. Moreover, nonparametric structural models of auctions are often just identified, as we can often perfectly rationalize observed bids with an appropriately constructed model (see Elyakime, Laffont, Loisel and Vuong (1994), Guerre, Perrigne, Vuong (2000) and Campo, Guerre, Perrigne, Vuong (2000)). To some skeptical researchers, the fact that the theoretical model has few testable implications may limit the usefulness of structural auction models.

In this paper, our goal is to assess whether structural models of first-price auctions can generate reason-

---

[3] Hendricks, Pinkse and Porter (2003) conduct a similar test.

[4] See, for example, Bajari and Ye (2003) in the first-price procurement auction context or Hortaçsu and Puller (2004) in multi-unit auctions for electricity generation. Related papers by Genesove and Mullin (1998) and Wolfram (1999) compare markups estimated using Cournot competition assumptions to markups obtained using cost data.

able estimates of bidders' private information. We structurally estimate first-price auction models using data from laboratory experiments. An advantage of laboratory experiments, compared to field data, is that bidder valuations are known. Therefore, we can directly compare our estimates of the structural parameters with valuations assigned to bidders in the experiment. This exercise will allow us to directly test whether a structural model can correctly recover bidder valuations at least in an experimental setting. Observing both the bids and the valuations leads to many more overidentifying restrictions and much greater power when testing the theory.

We structurally estimate four alternative first-price auction models using bidding 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) an adaptive model of learning, and 4) Quantal Response Equilibrium (QRE). 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 estimating the risk neutral Bayes-Nash model using the nonparametric methods suggested by Elyakime, Laffont, Loisel and Vuong (1994) and Guerre, Perrigne, Vuong (2000). This method involves two steps. First, the economist nonparametrically estimates the distribution of bids. Second, the economist uses a bidder's first order conditions to recover the valuation. Since the valuations are controlled for in the experiment, we can compare structural estimates of the bidder valuations to the true valuations. Because these methods are nonparametric, prior knowledge of the parametric family of the distribution of valuations is not required for this comparison.

Next, we nonparametrically estimate a model in which bidders are risk averse. We consider the symmetric constant relative risk aversion (CRRA) model of Holt (1980), Harris and Raviv (1980) and Riley and Samuelson (1981). In first-price auction experiments, observed bids are often found to be greater than the Nash equilibrium bids. Cox, Smith and Walker (1983a,b;1985a,b;1988) suggest that risk aversion is a potential explanation to "overbidding" in first-price auction experiments. The method we use builds on a technique proposed by Campo, Guerre, Perrigne, Vuong (2000).

The first two models we examine are all fully rational models that impose Bayesian Nash equilibrium behavior. Since many economists are skeptical of the rationality assumptions used in structural modeling, we also estimate two models where bidders are not perfectly rational. Our first "less-than-rational" model allows for adaptive learning. In this model, bidders do not have rational expectations regarding their competitors' bidding behavior. Instead, we assume that bidders beliefs about the distribution of bids are formed based on previous rounds of the experiment. Bidders then maximize expected utility given these beliefs. We assume that bidders' beliefs are formed using standard tools for nonparametric estimation of distributions. This approach is inspired by Sargent (1993), who considers models where agents form beliefs like an econometrician forms beliefs. Similar approaches have been used in Bray (1982), Marcet and Sargent (1989) and Cho and Sargent (1997). We implement the Sargent approach since it is the only approach we are aware of that generates an econometrically tractable specification for learning in first-price auction games.[5]

The second "less-than-rational" model is McKelvey and Palfrey's Quantal Response Equilibrium (QRE).

---

[5] 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 tends 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.

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 popularity 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.[6] Goeree, Holt, and Palfrey (2002) apply 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. We develop a simple maximum likelihood procedure to estimate a symmetric QRE model with risk aversion after making parametric assumptions on the distribution of valuations.

We find that the symmetric CRRA Bayes-Nash model is able to recover the distribution of valuations better than the risk neutral Bayes-Nash model and adaptive learning model. In fact, a statistical test based on a modified Kolmogorov-Smirnov test statistic fails to reject the equality of estimated and actual distribution of bidder distributions under the CRRA specification. The QRE model with risk aversion uncovers a very similar risk aversion parameter as the Bayes-Nash model, though the CRRA Bayes-Nash model yields more accurate estimates of the underlying distribution of bidder valuations.

This paper makes three main contributions to the empirical literature on auctions. First, to the best of our knowledge, previous work in the structural econometric literature has not compared the estimated valuations to the true valuations in order to assess the ability of structural first-price auction models to recover the distribution of bidder valuations.[7] Second, we illustrate some potential strengths and weakness

---

[6] See Goeree and Holt (1999) and Capra, Gomez, Goeree, and Holt (2002) among a large and growing literature.

[7] In independent work, Plott and Salmon (2004) have structurally estimated a model of bidder behavior in the simultaneous ascending auction using experimental data. They then compare the estimated and actual valuations to assess the plausibility of the method. A distinction between our work and theirs is that their structural econometric model is not based on a equilibrium model of bidding, but on "heuristic" behavioral restrictions.

of various 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. Finally, our paper suggests a new application of experimental economics. We use experiments to assess the merits of competing structural models. While our application concerns auction models, experimental data could be used in an analogous fashion to study other types of structural models.

Several words of caution are in order when interpreting our results. First, experimental environments may differ significantly from "real" economic environments. In real auctions, the stakes 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 Bayesian-Nash equilibrium model with risk aversion 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.

2 The Data.

The data set was provided to us by John Kagel and contains results from first-price auction experiments conducted by Dyer, Kagel, and Levin (1989). The subjects were primarily recruited from MBA students at the University of Houston. There were three experimental runs with six different subjects participating in each run for a total of 18 subjects. In these experiments, bidders were assigned i.i.d. valuations v drawn from a uniform distribution on $[\$0, \$30]$. In the event that they won, subjects were paid their valuation minus their bid. Each subject participated in 28 auctions over the course of two hours. As in the analysis of Dyer, Kagel, and Levin, we exclude data from the first five runs of the experiments. This leaves us with three runs of 23 auctions.

The number of bidders was determined at random in the experiment. With probability $1/2$, there were $N = 3$ bidders and with probability $1/2$, $N = 6$ bidders. Subjects submitted two "contingent" bids and one "non-contingent" bid. After the bids were submitted, a coin was tossed 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 $N = 6$. Otherwise, the non-contingent bid was used so that the bid could not be conditioned on N. After each auction, bids and corresponding private values were publicly posted on a blackboard.

Throughout the rest of this paper, we will ignore the non-contingent bids and focus on the contingent bids.[8] The main advantage of using the Dyer, Kagel, and Levin (1989) data is that there is variation in the number of bidders. This variation will allow us to identify a broader class of economic models than if N is held fixed. Also, we will be able to explore the merits of alternative estimators as the number of bidders changes. We acknowledge that this is a non-standard experimental setup, since bidders are supposed to

---

[8] 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.

make three simultaneous bidding decisions instead of one. This may change their response and perhaps increase the frequency of mistakes. However, as we will discuss 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.

In this game, if bidders are risk neutral, equilibrium bidding strategies in this symmetric, independent private values model will be given by:

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

In equation (1), v is a bidder's private valuation for winning the auction and $\mathbf{b}(v)$ is the equilibrium bid function. The bid functions predicted by equilibrium are linear, with a slope of $\frac{2}{3}$ when there are $N = 3$ bidders and $\frac{5}{6}$ when $N = 6$.

The observed bids differ considerably from the bids predicted by the risk neutral Bayes-Nash equilibrium. Figure 1 plots the observed bids and the equilibrium bids in the $N = 3$ and $N = 6$ bidder auctions. Clearly, the observed bids are higher than the bids predicted by the Bayes-Nash equilibrium, particularly when there are $N = 3$ bidders. This has been referred to as the "overbidding" phenomenon. Overbidding is commonly observed in first price auction experiments see, for instance, Cox, Smith, and Walker (1988), Harrison (1989) and Kagel (1995). In table 1, we summarize the difference between the equilibrium and actual bids. The mean observed bid is $\$2.44$ higher than the Nash bid when there are three bidders, about 16% of the mean valuation. When $N = 6$, the observed bids are higher, but this time by the smaller amount of $\$0.65$. Therefore, overbidding is most pronounced when there are three bidders.

Next, we ask how close bidders are to maximizing expected profits in the auction. We begin by estimating the bidder expected profits in the experiment. Let $Q(b; N, e)$ denote the cumulative distribution of bids with N bidders in experiment e. If a bidder has a valuation of v, then her expected profit from submitting a contingent bid of b in an auction with N bidders is

$$\pi(b, v; N, e) = (v - b) \cdot Q(b; N, e)^{N-1} \tag{2}$$

Let $\widehat{Q}(b; N, e)$ denote an estimate of $Q(b; N, e)$. Then our estimate of the expected profit for submitting a contingent bid of b will be:

$$\widehat{\pi}(b, v; N, e) = (v - b) \cdot \widehat{Q}(b; N, e)^{N-1} \tag{3}$$

If a bidder is risk neutral, we define the optimization error as the difference between the profit maximizing bid and the bid actually submitted. If a bidder submits a contingent bid of b, the optimization error $\omega(b, v; N, e)$ is estimated as:

$$\omega(b, v; N, e) = \left(\arg\max_{b'} (v - b') \cdot \widehat{Q}(b'; N, e)^{N-1}\right) - \widehat{\pi}(b, v; N, e). \tag{4}$$

In equation (4), the first term, $\left(\arg\max_{b'} (v - b') \cdot \widehat{Q}(b'; N, e)^{N-1}\right)$ is the utility maximizing bid given an estimate $\widehat{Q}$ of the distribution of bids. The second term, $\widehat{\pi}(b, v; N, e)$ is the utility that the bidder received from the bid observed in the experiment. We estimated $\widehat{Q}$ by smoothing the empirical cdf using a normal kernel. The bandwidth was selected using Silverman's rule-of-thumb.[9] Since bidder valuations

---

[9] We could also estimate the distribution of valuations by computing the optimal bandwidth. However, when estimating the learning model, we will sometimes have to estimate the distribution of bids based on a small number of observations. In such cases, estimating the distribution of bids with an optimal bandwidth makes little sense. We present our results using Silverman's rule of thumb in order to have a consistent approach for estimating the distribution of bids. With the exception of the learning model, this will have very little influence on our results.

are contained in our data set, we can estimate expected profits, $\widehat{\pi}(b, v; N, e)$ and the optimization error $\omega(b, v; N, e)$ by applying equations (3) and (4).

We summarize these results also in table 1. The expected payoffs in this experiment are fairly modest. Conditional on $N = 3$, bidders have an expected payoff of \$1.33 and \$0.58 when $N = 6$. As a consequence, in monetary terms, the optimization error that bidders make is also fairly small. The median optimization error is \$0.19 in the three bidder auctions and less than one cent in the six bidder auctions. Therefore, while the bids may differ substantially from equilibrium bids, this results in only small monetary losses for the participants.[10]

In summary, our description of the data contains mixed results about the success of the theory. The equilibrium bids are quite different from those predicted by standard theory, especially in the three bidder auctions. However, the main assumption of the theory is that bidders maximize expected utility. This condition appears to be approximately satisfied. By any reasonable metric, the losses that occur from submitting suboptimal bids are fairly small. A priori, it is not clear how these deviations from equilibrium will influence our structural estimates. Is the fact that bidders are close to optimizing sufficient to get reasonable estimates of valuations? Will models that modify the benchmark, risk-neutral model generate superior estimates of the structural parameters?

In the following sections, we will structurally estimate four alternative models of the first-price sealed-bid auction. We first briefly describe each model. Next, identification and estimation of the model is discussed. Finally, we compare the estimated valuations with the valuations assigned to bidders in the experiments to determine whether the model generates reasonable estimates.

---

[10] These findings are consistent with Harrison (1989).

3 The Risk Neutral Model.

We begin by considering the benchmark risk neutral model of bidding in the first-price sealed-bid auction. There are $i = 1, ..., N$ symmetric bidders. Bidder i's valuations v_i is private information that is independently and identically distributed with cdf $F(v)$ and pdf $f(v)$. Bidders simultaneously submit sealed bids b_i. If bidder i's bid is the highest, her utility is $v_i - b_i$, and is zero otherwise.[11]

Let $b = \mathbf{b}(v)$ denote the equilibrium bid function. Under weak regularity conditions, the equilibrium bid function is strictly increasing and differentiable so that its inverse $\phi(b)$ exists and inherits these properties. Bidder i's profit from bidding b_i is:

$$\pi_i(b_i; v_i) \equiv (v_i - b_i) F(\phi(b_i))^{N-1}. \tag{5}$$

In equation (5), 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, $F(\phi(b_i))^{N-1}$.

3.1 Identification and Estimation

The first order condition for maximizing expected utility can be written as

$$-F(\phi(b_i))^{N-1} + (N-1)(v_i - b_i) F(\phi(b_i))^{N-2} f(\phi(b_i)) \phi'(b_i) = 0 \tag{6}$$

$$v_i = b_i + \frac{F(\phi(b_i))}{f(\phi(b_i)) \phi'(b_i)(N-1)} . \tag{7}$$

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 (7) as

---

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

$$v_i = b_i + \frac{G(b_i)}{g(b_i)(N-1)} \,. \tag{8}$$

An empirical implication of equation (7), first exploited by Elyakime, Laffont, Loisel and Vuong (1994), and expanded upon in Guerre, Perrigne and Vuong (2000), is that equation (8) suggests a simple estimator. Suppose that the econometrician observed T repetitions of the auction. Let $b_{i,t}$ denote the bid that i submits in auction t. Since we have multiple repetitions of the same auction, it is possible to estimate G and g. Denote these estimates as $\hat{G}(b)$ and $\hat{g}(b)$. If we substitute the estimated distributions into equation (8), we can generate an estimate of $\widehat{v}_{i,t}$ of $v_{i,t}$, bidder i's valuation in the tth auction as follows:

$$\widehat{v}_{i,t} = b_{i,t} + \frac{\widehat{G}(b_{i,t})}{\widehat{g}(b_{i,t})(N-1)}. \tag{9}$$

Guerre, Perrigne and Vuong (2000) show that we can consistently estimate the distribution and density of the pseudo-valuations, $\widehat{v}_{i,t}$, using nonparametric methods, except within a neighborhood of the boundaries of its support.[12] By applying equation (9) to every bid in our data set, we can generate estimates, $\{\widehat{v}_{i,t}\}_{i=t,..,N,\, t=1,...,T}$ of the valuations associated with each bid in our data set.

To summarize, the estimation procedure involves two steps:

1. First, using nonparametric methods generate estimates $\widehat{G}$ and $\widehat{g}$ of G and g.
2. Given the first stage estimates, apply equation (9) for every observed bid $b_{i,t}$ to generate $\widehat{v}_{i,t}$, an estimate of $v_{i,t}$.

For a detailed discussion of the asymptotic properties of this estimator, the interested reader is referred to Guerre, Perrigne and Vuong (2000). Versions of this estimator are also considered by Elyakime, Laffont, Loisel and Vuong (1994), Guerre, Perrigne, Vuong (2000), Campo, Guerre, Perrigne and Vuong (2003), Flambard and Perrigne (2002), Bajari and Ye (2003) and Jofre-Bonet and Pesendorfer (2003).

---

[12] Guerre, Perrigne and Vuong (2000) suggest trimming the sample of bids near its boundaries.

It is easy to see that each one of the valuations $v_{i,t}$ is just identified. Given a bid $b_{i,t}$, all of the terms on the right hand size of the equation (9) are known. By associating a valuation $\widehat{v}_{i,t}$ with each bid $b_{i,t}$, we can perfectly rationalize all of the observed bids in the auction. In principle, we could test whether the relationship between $\widehat{v}_{i,t}$ and $b_{i,t}$ is strictly monotonic. If this failed, we could reject that the bid functions are increasing, an overidentifying restriction from the theory. However, if this monotonicity condition is satisfied, we will be able to perfectly rationalize the observed bids with our model. By only observing the bids, assuming that the data generating process is a risk-neutral Bayes-Nash equilibrium model is an identifying assumption that it is difficult to either verify or refute.

3.2 Assessing Goodness of Fit

Next, we need to assess the goodness of fit of our model by comparing the estimated distribution of valuations with the actual (assigned) distribution of valuations. The first metric we use to compare these two distributions is a modified Kolmogorov-Smirnov (KS) statistic. As is well known, the KS test statistic for the equality of the empirical distribution $\hat{F}_T(v)$ (based on an i.i.d. sample with T realizations) and the (known) true distribution, $F(v)$, on the support $[\underline{v}, \bar{v}]$ is given by the normalized distance:

$$KS_T = \sqrt{T} \sup_{v \in [\underline{v}, \bar{v}]} |\hat{F}_T(v) - F(v)|$$

In our application, however, the empirical distribution function, $\hat{F}_T(v)$, is not based on i.i.d. realizations from $F(v)$, but on a sample of pseudo-valuations $\{\widehat{v}_{i,t}\}_{i=t,..,N,\, t=1,...,T}$, estimated using the relation:

$$\hat{v}_{i,t} = b_{i,t} + \frac{1}{N-1} \frac{\hat{G}(b_{i,t})}{\hat{g}(b_{i,t})} \tag{10}$$

Hence, our test-statistic needs to take into account the sampling error associated with the first step estimates of $\hat{G}$ and $\widehat{g}$. Therefore, we study the asymptotic distribution of the modified KS statistic:

$$MKS_T = \sqrt{T} \sup_{v \in [\underline{v}, \bar{v}]} |\frac{1}{T} \sum_{t=1}^{T} 1\{\hat{v}_t \leq v\} - F(v)|$$

with $\hat{v}_t$ given as above, where $1\{.\}$ is the indicator function.

A first intuition might be to bootstrap MKS_T to take into account the first-step sampling error. However, it is not obvious whether this test statistic has an asymptotic distribution that is normal, which is the necessary and sufficient condition for bootstrap to work according to Mammen's theorem (Horowitz (2001)). Following an application in Haile, Hong and Shum (2003), a subsampling strategy, instead of the bootstrap, may be used in this case, since subsampling is consistent for test statistics with general non-degenerate asymptotic distributions.[13]

To do this, we will modify MKS_T to allow for a smooth analogue for the estimator of the empirical distribution function:

$$M\tilde{K}S_T = \sqrt{T} \sup_{v \in [\underline{v}, \bar{v}]} |\frac{1}{T} \sum_{t=1}^{T} \Lambda(\hat{v}_t - v) - F(v)|$$

where we take $\Lambda(x)$ to be $1 - \psi(\frac{x}{h'})$, where $\psi(.)$ is a smooth, strictly monotonic distribution function and h' is a bandwidth parameter. As $h' \to 0$, $M\tilde{K}S_T \to MKS_T$.

Under the null hypothesis of risk neutrality (the risk aversion case is similar), Guerre, Perrigne and Vuong (2000) prove the uniform consistency of the nonparametric estimator (10) over a support $[\underline{v}, \bar{v}]$ strictly bounded away from the upper and lower supports of the valuation distribution. Thus, since Λ is strictly monotonic, $M\tilde{K}S_T \to 0$ under the null hypothesis.

---

[13] Haile, Hong and Shum (2003) use a subsampling approach to construct a statistic to test for the null of independent private values in first-price auctions. Their test can be thought of as a one-sided test for first-order stochastic dominance, and is for a very different application than is considered here. However, their main testing strategy applies in this context.

The subsampling algorithm is implemented as follows. Let R_T be a sequence of subsample sizes, and let

$$\kappa_T = \binom{T}{R_T}$$

be the number of unique subsamples in a data sample of size T. The sampling distribution Φ_T of the test statistic $M\tilde{K}S_T$ is approximated by:

$$\Phi_{T\,,R_T}(x) = \frac{1}{\kappa_T}\sum_{i=1}^{\kappa_T} 1\{\sqrt{R_T}\sup_{v\in[\underline{v},\bar{v}]}|\frac{1}{R_T}\sum_{t=1}^{R_T}\Lambda(\hat{v}_t - v) - F(v)| \le x\}$$

If $M\tilde{K}S_T$ possesses a nondegenerate limiting distribution, then Theorem 2.6.1 in Politis, Romano, and Wolf (1999) shows that this approximation will yield a consistent estimate of the true sampling distribution under the null, if $R_T \to \infty$ and $\frac{R_T}{T} \to 0$ as $T \to \infty$. In the appendix, we derive the limiting distribution of $M\tilde{K}S_T$ to establish that the use of subsampling is justified.

We also assess the goodness of fit of our model by comparing the distance between the estimated and actual valuations in the L^1 and L^2 norms, defined as:

$$\begin{aligned} L^1 &= \frac{1}{T}\sum_t |\widehat{v}_t - v_t| \\ L^2 &= \left(\frac{1}{T}\sum_t (\widehat{v}_t - v_t)\right)^{1/2} \end{aligned}$$

Although we have not derived the statistical properties of these distance metrics, and hence can not conduct formal hypothesis tests based on them, these metrics have some intuitive content (mean absolute and mean square deviation, measured in units of dollars), and provide an alternative means to compare the performance of the various models we will employ.

3.3 Results.

To estimate the symmetric risk neutral Bayesian-Nash model, the strategy described in section 2.1 is followed. When estimating the density of bids, $g(b)$, we use the normal kernel with Silverman's rule-of-thumb bandwidth. The cdf $G(b)$ is estimated using the empirical cdf. We pooled all bids across all the experiments when estimating these densities.[14] To recover $\widehat{v}_{i,t}$, the valuation associated with the bid $b_{i,t}$, equation (9) is used.

In figure 2, we plot the histograms of estimated valuations for the $N = 3$ and $N = 6$ cases. The actual valuations were assigned in the experiment and were distributed uniformly on $[0, 30]$. The valuations estimated from the three bidder auctions are visually quite different from the distribution of the actual valuations, particularly on the upper tail of the distribution. A fairly significant fraction of the valuations are greater than 30. In the six bidder auction, however, the estimator appears to do a much better job. This is not surprising given our results about overbidding. In the three bidder auctions, overbidding was much more pronounced than in the six bidder auction.

To implement the modified KS testing procedure described in the previous section, we evaluated the test statistic at random subsamples, since it is not feasible to enumerate all possible $\binom{T}{R_T}$ subsamples. We used 500 subsamples, each of size $R_T = 200$. For each subsample, we set the bandwidth of the kernel density estimators used in $\hat{g}(b_t)$ to be $h_{R_T} = O(R_T^{-1/5})$ using Silverman's rule-of-thumb. We set $h' = 1$ but experimenting with $h' = 2$ and $h' = 0.1$ did not yield different results.

Table 2 reports the results of the structural estimation and testing procedure. For the $N = 3$ case, the

---

[14] We estimated g and G both by pooling observations across experiments and by estimating them separately for each experiment. In the results we report, we opt for the later specification. This makes little difference for the symmetric model, but is more consistent with how we will estimate the risk averse and learning model.

p-value of the modified KS statistic was essentially zero for both trimming levels (at the 25-75th and 5-95th percentiles of the valuation distribution), i.e. we were able to reject the equality of the estimated valuation distribution with the true distribution. For the $N = 6$ case, we were not able to reject the equality of the estimated and actual distributions at the 5% confidence level.

We also report the results for the L^1 and L^2 norms in Table 2. The three bidder model generates poor results in both norms. The expected and actual valuation have on average a difference of nearly $\$4$. In the $N = 6$ case, however, the estimated and actual valuations are, on average, within $\$1.07$. To summarize, the estimates are quite reasonable when the number of bidders is equal to six, but are quite poor when $N = 3$. These results are generated by the sizeable overbidding in the experiments with three bidders. These results suggest that we should consider alternatives to the standard risk neutral model, particularly when the number of bidders is small.

4 The Risk-Averse Model.

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 (1983a,b;1985a,b;1988) note that one possible explanation for overbidding is risk aversion, and offer a model with risk averse bidders to explain experimental results in first-price auctions. Indeed, the data set of Dyer, Kagel and Levin (1989) was used to test a comparative static implication of equilibrium bidding with risk aversion and found support for this hypothesis.

In light of the experimental literature, we add risk aversion in a parsimonious manner to our econometric specification. In particular, we assume that bidders have a constant-relative-risk-aversion (CRRA) utility function, $U(x) = x^\theta$, $\theta \in [0, 1]$. In this specification, $1 - \theta$ is the coefficient of relative risk aversion, with

$\theta = 1$ corresponding to risk neutrality.

In this model, the first order condition is:

$$v_i = b_i + \theta \cdot \frac{G(b_i)}{g(b_i)(N-1)}. \tag{11}$$

Observe that when bidders are risk neutral, that is $\theta = 1$, then (11) reduces to (7).[15]

4.1 Structural Estimation

The logic of the estimator is similar to the previous section. If the economist knew G and θ, then we could construct a two-step estimator along the lines of the previous section. The problem that we face, however, is that θ is not directly observed. Therefore, we must find a way to estimate it from the data.

We use a method based on the techniques proposed in Campo, Guerre, Perrigne and Vuong (2000). Let $G(b; N)$ denote the distribution of bids with N bidders. Let v_α denote the α^{th} percentile of the distribution of valuations. Let $b_\alpha^{(3)}$ denote the α^{th} percentile of $G(b; 3)$ and let $b_\alpha^{(6)}$ denote the α^{th} percentile of $G(b; 6)$. By equation (11) it follows that

$$v_\alpha = b_\alpha^{(3)} + \theta \cdot \frac{G(b_\alpha^{(3)}; 3)}{2g(b_\alpha^{(3)}; 3)} \tag{12}$$

$$v_\alpha = b_\alpha^{(6)} + \theta \cdot \frac{G(b_\alpha^{(6)}; 6)}{5g(b_\alpha^{(6)}; 6)} \tag{13}$$

By simple algebra, it follows from the equations (12) and (13) that:

---

[15] Note also that for uniformly distribution valuations, equilibrium bid functions are given by:

$$b(v_i) = \frac{N-1}{N-1+\theta} v_i$$

as derived by Holt (1980), Riley and Samuelson (1981), Harris and Raviv (1981), and that for $\theta > 0$ this predicts that risk aversion leads to more aggressive bidding than in the risk neutral case.

$$b_\alpha^{(3)} - b_\alpha^{(6)} = \theta \cdot \left(\frac{G(b_\alpha^{(6)}; 6)}{5g(b_\alpha^{(6)}; 6)} - \frac{G(b_\alpha^{(3)}; 3)}{2g(b_\alpha^{(3)}; 3)} \right) \tag{14}$$

Equation (14) suggests a simple way to estimate θ. If we knew the distribution of bids in the three and six bidder experiments, given α, all of the terms on the left and right hand in this equation would be directly observable except for θ. By evaluating (14) at a large number of percentiles, we could then estimate θ using regression. Given an estimate $\widehat{\theta}$ of θ, we can then estimate the valuations v_i by evaluating the empirical analogue of equation (11) as in the previous section.

To summarize, we generate estimates $\widehat{v}_{i,t}$ of $v_{i,t}$ as follows:

1. Generate nonparametrically estimates $\widehat{G}(b; N)$ and $\widehat{g}(b; N)$ of $g(b; N, e)$ and $G(b; N, e)$.
2. Generate an estimate $\widehat{\theta}$ of θ by running the following regression, using a finite number of percentiles α:

$$\widehat{b}_\alpha^{(3)} - \widehat{b}_\alpha^{(6)} = \theta \cdot \left(\frac{\widehat{G}(\widehat{b}_\alpha^{(6)}; 6)}{5\widehat{g}(\widehat{b}_\alpha^{(6)}; 6)} - \frac{\widehat{G}(\widehat{b}_\alpha^{(3)}; 3)}{2\widehat{g}(\widehat{b}_\alpha^{(3)}; 3)} \right) + \varepsilon_\alpha \tag{15}$$

3.Given $\widehat{\theta}$, $\widehat{G}(b; N)$ and $\widehat{g}(b; N)$ use the empirical analogue of (11) to generate an estimate $\widehat{v}_{i,t}$ of $v_{i,t}$.

$$\widehat{v}_{i,t} = b_{i,t} + \widehat{\theta} \cdot \frac{\widehat{G}(b_{i,t})}{\widehat{g}(b_{i,t})(N-1)} \tag{16}$$

The results of Campo, Guerre, Perrigne and Vuong (2000) demonstrate that it is possible, in principle, to recover θ using data from just a single percentile. Since our concern is with testing, in order to improve the efficiency of our estimator, in step 2 we use data from a large number of percentiles when estimating θ. Otherwise, the logic of our estimator is analogous to Campo et al. (2000).

Equation (11) suggests that given a value of θ, bidder valuations are just identified as in the risk neutral model. We can associate each bid with a valuation and therefore perfectly rationalize what we see in the data with valuations constructed as in equation (16). The estimation procedure makes an identifying assumption that the distribution of private information is identical when $N = 3$ and $N = 6$. This assumption was stronger than necessary in order to estimate the risk neutral model of the previous section. Given this stronger assumption, we can make stronger conclusions about the data generating process. In particular, we will be able to nest the risk neutral model as a special case (i.e. $\theta = 1$). Also, we can in principle test whether other restrictions of the theory such as whether the distribution of valuations is in fact equal when $N = 3$ and $N = 6$ by applying (16) separately for these two cases. However, from the bids, we will not be able to verify that the data generating process is the symmetric risk averse model. Moreover, even if we reject the model, we cannot be certain that our rejection is based on small deviations from expected utility maximization as we found in section 2. Therefore, it is useful to compare actual and predicted valuations and one mechanism to assess the performance of this model.

4.2 Estimates

We estimated the symmetric risk aversion model using a variety of specifications. The estimates of G and g were the same as in the previous section. The main factor that appeared to affect risk aversion coefficient estimates was how we trimmed the boundaries of the bid distribution. As reported in table 3, the estimated θ is about 0.16 when the entire data is used, but goes up to 0.22-0.25 when we trim the upper and lower boundary of the data. Visual inspection of the independent and dependent variables in the regression in equation (14) revealed a monotonic relationship when the upper boundary of the support was trimmed at 95% of bids. For higher bids, the association between dependent and independent variables appeared to be negative, rather than positive, suggesting that a simple linear relationship afforded by the CRRA specification can not account for the highest bids in the data set.

Note that an additional verification of the risk aversion hypothesis can be obtained by utilizing our knowledge of the experimental setup. As noted above, the Bayesian Nash equilibrium bid functions with symmetric CRRA risk aversion $\alpha = 1 - \theta$ and uniformly distributed valuations are given by the formula

$$b(v) = \frac{N-1}{N-\alpha} v$$

Since we have data from both $N = 3$ and $N = 6$ bids, if we impose the equality of the risk aversion coefficient α across the two treatments, we get:

$$\frac{b^{(3)}(v)}{b^{(6)}(v)} = \frac{2}{5} \frac{6-\alpha}{3-\alpha}$$

where $b^{(3)}(v)$ is the $N = 3$ bid of a bidder with valuation v, and $b^{(6)}(v)$ is the $N = 6$ bid. We can thus get an estimate of α by estimating the mean of the ratio of $N = 3$ and $N = 6$ bids, and solving for α. When we implemented this in the data, we found $\theta = 1 - \alpha = 0.2259$.

We then used our point estimate of the risk aversion coefficient to estimate the individual valuations. We use the same three distance measures (modified KS-distance, L^1 and L^2 norms) to compare the estimated valuations with the actual valuations. Once again, we use the subsampling approach to approximate the distribution of the KS-statistic, with the additional complication of recomputing the risk aversion coefficient for each random subsample of the data. Since the risk aversion coefficient estimate uses data from both $N = 3$ bids and $N = 6$ bids, and the test statistics across the two cases are not independent, we also considered a “joint-test statistic” which is the sum of the KS statistic for the $N = 3$ and $N = 6$ cases. Once again, the distribution of this statistic was approximated by subsampling.

Table 4 reports the results of the testing procedure. When we trim the data to both the 25-75th and 5-95th percentiles of the bid and valuation supports, the values of the test statistics are lower than those reported in table 2, and we fail to reject the equality of the estimated and actual valuation distributions for both $N = 3$ and $N = 6$. However, the performance of the estimator deteriorates somewhat when we begin to include the highest bids and valuations in the data. Interestingly, in contrast to table 2, the performance of the estimation routine is better for the $N = 3$ case, whereas the estimates are somewhat worse for the $N = 6$. Apparently, the risk averse specification attempts to compensate for the observed overbidding in the $N = 3$ case, but, in doing so, begins to underestimate the valuations rationalizing the $N = 6$.

In table 4, we also report the L^1 and L^2 norms (again using the risk aversion coefficient estimated with the 25-75th percentile of the data). The average absolute difference was approximately \$1.35 in both the $N = 3$ and $N = 6$ case. The median absolute difference was only \$0.85 for the $N = 3$ case and \$1.04 in the $N = 6$ case.

In figure 3, we plot the distribution of the actual valuations and the estimated valuations corresponding to the $N = 3$ and $N = 6$ bidder cases, where the risk aversion coefficient is estimated using the 25-75th percentiles of the data . The estimated distribution of valuations appears to be uniform, except on the right tail for valuations greater than 25. In table 5, we display the percentiles of the actual and estimated distribution of valuations. The estimated and actual distribution of valuations agree quite closely under the 80^{th} percentile. The results of the KS test reflect this discrepancy in the tails of the distributions. As we can see in table 5, the percentiles of the estimates and actual distributions of valuations are quite close except for the right tail.

We thus conclude that the symmetric risk aversion specification does a much better job than the risk neutral specification of the previous section in the $N = 3$ case. When $N = 6$, the estimates are quite comparable. This result suggests that estimates of bidder valuations will be more sensitive to the choice of method when the number of bidders is small. Second, our estimation method exploits variation in the number of bidders assuming that the distribution of valuations remains constant across the two auctions. In order to implement such an estimator in the field, the economist would have to make some assumptions about the distribution of valuations across auctions. Also, the economist would have to control for the endogeneity of the number of bidders. Such assumptions are made explicitly (or implicitly) in a number of recent empirical studies.[16] Third, it is worth noting that variation in the number of bidders allows us to identify a more general model of preferences. Variation in N allows us to distinguish risk neutrality from risk aversion by running equation (15).

5 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 $b_{i,t}$. Formally, we define the history as follows. Fix an experiment, $e = 1, 2, 3$. Let $t = 1, ..., T$ denote the t^{th} repetition of the auction game in experiment e. The history is the set of all bids that the subject would have seen from previous plays of the game, i.e., $h_{it} = \{b_{i,r}\}_{r<t, i=1,...,N}$.

Just as the econometrician has to estimate $G(b)$ using the empirical distribution of bids, we assume in this model that bidders form beliefs about $G(b)$ using previously submitted bids. We denote these beliefs as $G(b|h_{it})$. Given their beliefs, bidders choose their bids in order to maximize expected profit $\pi_i(b_i; v_i, G(b|h_{it}))$ which is equal to

---

[16] See, for example, Haile, Hong and Shum (2003).

$$\pi_i(b_i; v_i, G(b|h_t)) = (v_i - b_i)G(b|h_{it})^{N-1}. \tag{17}$$

In the experiment we consider, bidders were told after each auction what their opponents' bids were, so that $h_{it} = h_t$, all bids submitted until t.

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

$$v_i = b_i + \frac{G(b_i|h_t)}{g(b_i|h_t)(N-1)} \; . \tag{18}$$

In order to implement this estimator, we must model agent's beliefs conditional on h_t. There now exists a fairly substantial theoretical literature on learning in games which models agents as making best responses to their beliefs which depend on past plays of the game. Unfortunately, most of these specifications do not generate econometrically tractable models. We choose to follow the modeling approach of Sargent (1993) and assume that agents form beliefs like econometricians. We will assume that G and g are formed by using the nonparametric methods that we described in the previous sections. We will let $\widehat{G}(b|h_t)$ and $\widehat{g}(b|h_t)$ denote our estimates of the agent's beliefs. Note that the estimates of the bidders' beliefs only depend on the plays of the game that they have viewed.[17]

The model 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

---

[17] In practice, it does not make sense to estimate the density $\widehat{g}(b|h_t)$ by using the optimal bandwidth, since this will be imprecisely estimated during early rounds. Therefore, we assume a normal kernel and use Silverman's rule of thumb. The bidders would not be able to compute the sample standard deviation conditional on h_t (except for the last time period t). However, this choice seemed desirable assuming that a priori the bidders knew the parametric family of $\widehat{g}(b|h_t)$.

actors, and where previous bids are publicly observable.[18] However, there is mixed empirical evidence that having better data on past bids allows for better, or different bidding decisions.[19]

As in section 3, we can estimate $\hat{v}_{it}$ using a two stage procedure. We first generate estimates $\widehat{G}(b|h_t)$ and $\widehat{g}(b|h_t)$ of $G(b|h_t)$ and $g(b|h_t)$. Then we generate an estimate $\hat{v}_{it}$ of v_{it} by using the empirical analogue of (18), i.e.

$$\widehat{v}_{i,t} = \widehat{b}_{i,t} + \frac{\widehat{G}(\widehat{b}_{i,t}|h_t)}{\widehat{g}(\widehat{b}_{i,t}|h_t)(N-1)} \,. \tag{19}$$

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.

As with the risk neutral model, valuations in the learning model are just identified. Each $\widehat{b}_{i,t}$ can be associated with a valuation, $\widehat{v}_{i,t}$, that perfectly rationalizes this bid. It is not clear that the theory generates any overidentifying restrictions. For example, if agents are learning, it is not obvious that bid functions should be monotonic under standard valuation distributions. Based on observations of bids alone, it seems very difficult to possibly refute the learning model. Therefore, once again it is useful to compare the estimated and actual valuations as a means to assess the usefulness of the learning model.

---

[18] Such as many procurement auctions. 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.

[19] Most empirical studies on this issue focus on whether more experienced bidders make better, or at least, 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. Ockenfels and Roth (2003) and Bajari and Hortaçsu (2003) 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 Hortaçsu (2003), 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 end up enjoying higher profits.

5.1 Estimates From the Learning Model.

In figure 4, we compare the estimated and actual valuations from the learning model. The results from the learning model appear to mirror the results from the risk neutral model. When $N = 3$, the estimated distribution of valuations differs considerably from the true distribution of valuations on the right tail. As in the estimates of the risk neutral model, we believe this is due to overbidding. We interpret these results as suggesting that learning cannot explain overbidding behavior and hence, it does not generate improved estimates of the structural parameters. In this $N = 6$ case, however, the estimates appear to be much more reasonable.

Table 6 reports the distance between the estimated and actual valuations. The values of the KS distance measure for the $N = 3$ and $N = 6$ cases are similar, but larger in magnitude than the risk neutral model.[20] In the L^1norm, the distance between the estimated and actual valuations is \$5.09 for $N = 3$ and \$1.46 for $N = 6$, again similar, but larger than the risk neutral model. These results suggest that our learning model does not lead to better structural estimates than the rational models in the norms that we consider.

We note that the results from the learning model will depend on the specification of $G(b_{i,t}|h_t)$ and $g(b_{i,t}|h_t)$. The specification we chose was based on the suggestion by Sargent (1993) to allow bidders to form beliefs like econometricians. This specification was chosen because it was the only econometrically tractable specification that could be applied to our problem. Also, there is little evidence about how bidders form beliefs in auctions. We believe that future research on how agents actually form beliefs in auctions might lead to more useful structural models that allow for learning.

---

[20] Unfortunately, it is difficult to conduct a hypothesis test using this distance measure, since assessing the sampling distribution of the statistic using subsampling or other resampling methods is very much complicated by the fact that the distribution and density of bids are estimated conditional on particular bid histories when constructing each $\widehat{v}_{i,t}$. The resampling procedure would have to take this conditioning into account in the proper manner.

6 The Quantal Response Equilibrium Model

Some recent research in experimental economics has made use McKelvey and Palfrey's (1995) Quantal Response Equilibrium (QRE) model to reconcile deviations from Nash equilibrium predictions. In section 2, our estimates suggested that bidders commonly failed to perfectly optimize and submitted bids that were quite different from those predicted by Bayes-Nash equilibrium. The QRE model is one way that experimental economists have chosen to model optimization error.

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, (5) and an idiosyncratic shock drawn from an i.i.d. extreme value (logit) distribution. 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. 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

$$\text{For all } v \in \mathcal{V} \text{ and } b \in \mathcal{B},\ \mathbf{B}(b|v) \geq 0. \tag{20}$$

$$\text{For all } v \in \mathcal{V},\ \sum_{b \in \mathcal{B}} \mathbf{B}(b|v) = 1. \tag{21}$$

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.

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:[21]

$$Q(b) = \left[\sum_{v} \sum_{b' < b} \mathbf{B}(b'|v) f(v) \right]^{N-1} . \tag{22}$$

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 $\varepsilon(b_i, v_i)$:

$$\widehat{\pi}(b; v_i) \equiv (v_i - b_i) * Q(b_i) + \varepsilon(b_i, v_i) = \pi(b_i; v_i) + \varepsilon(b_i, v_i) . \tag{23}$$

The logit equilibrium model generalizes the Bayes-Nash model by including the term $\varepsilon(b_i, v_i)$ in an agent's payoffs. In the experimental literature, the term $\varepsilon(b_i, v_i)$ is interpreted as the agent's optimization error. The optimization error $\varepsilon(b_i, v_i)$ is assumed to be i.i.d. 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 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).

---

[21] 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.

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:

$$\sigma(b_i; v_i, \mathbf{B}) = \frac{\exp(\lambda\pi(b_i; v_i, \mathbf{B}))}{\sum_{b' \in \mathcal{B}} \exp(\lambda\pi(b'; v_i, \mathbf{B}))}. \tag{24}$$

An equilibrium is a bidding function $\mathbf{B}(b|v)$ that is a fixed point of (24), that is $\mathbf{B}(b|v_i) = \sigma(b_i; v_i, \mathbf{B})$.

As $\lambda \to 0$, 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 $\lambda \to \infty$, 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 $\lambda \to \infty$. The existence of the logit equilibrium is obtained by using fixed point methods such as in McKelvey and Palfrey (1995).

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

$$p(b_i|\omega, \lambda; \widehat{Q}) = \int_{\underline{v}}^{\bar{v}} \sigma(b_i|v) f(v|\omega) dv = \int_{\underline{v}}^{\bar{v}} \frac{\exp(\lambda(v - b_i)\widehat{Q}(b_i))}{\sum_{b' \in \mathcal{B}} \exp(\lambda(v - b')\widehat{Q}(b'))} f(v|\omega) dv. \tag{25}$$

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

1. Given a data set of bids from T repetitions of the auction, $\{b_{i,t}\}_{i=1,\ldots,N,t=1,\ldots,T}$, form an estimate $\widehat{Q}(b)$ of $Q(b)$. In practice, we estimate $\widehat{Q}$ as the empirical c.d.f.
2. Estimate ω and λ using maximum likelihood, where the likelihood function for ω and λ is

$$L(\omega, \lambda) = \prod_{t=1}^{T} \prod_{i=1}^{N} p(b_{i,t} | \omega, \lambda; \widehat{Q}).$$

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

The above estimation procedure can be easily adapted to allow for risk aversion. For instance, if we use the CRRA specification of the previous section, equation (25) becomes:

$$p(b|\omega, \lambda, \theta; \widehat{Q}) = \int_{\underline{v}}^{\bar{v}} \sigma(b_i|v) f(v|\omega) dv = \int_{\underline{v}}^{\bar{v}} \frac{\exp(\lambda(v - b_i)^{\theta} \widehat{Q}(b_i))}{\sum_{b' \in \mathcal{B}} \exp(\lambda(v - b')^{\theta} \widehat{Q}(b'))} f(v|\omega) dv. \tag{26}$$

This might be a particularly desirable specification to take to the data, since Goeree, Holt, and Palfrey (2002) find that a QRE model with risk aversion provides a good fit to data from a first-price sealed-bid auction experiments. We should clarify once again, however, that in Goeree, Holt, and Palfrey (2002), the authors treated the valuations v as data and estimated λ and θ. In our "structural econometric estimation" exercise, we also need to estimate ω, the parameters characterizing the distribution of private information. The estimation task is much more demanding, and therefore, one should not expect this specification to be successful a priori. This is especially true since short of obtaining a global maximum of the likelihood function, we have not been able to obtain a formal identification result for this model. In fact, the recent work of Haile, Hortaçsu and Kosenok (2003) suggests that nonparametric identification of the QRE specification may not be possible in this setting if one abandons the iid assumption and allows for enough flexibility in the distribution of the idiosyncratic shock term.[22]

6.2 Results

Next, the valuation distribution is estimated assuming bidders are playing a symmetric logit (QRE) equilibrium. In forming the likelihood function, we assume that the econometrician knows the parametric form of the valuation distribution, though we recognize that in many real applications, most econometricians do not have this type of a priori information about the structural parameters.

We computed parameter estimates using three specifications. In the first specification, we assume that the econometrician knows that the distribution of valuations is uniform, but with unknown support given by the parameters $[vlower, vupper]$. We also impose the risk neutrality of the bidders. The second specification relaxes the risk-neutrality of the bidders by allowing the CRRA parameter θ to be different than 1, but retains the uniform distribution assumption. The third specification relaxes the uniform distribution assumption by assuming a beta distribution, with two additional parameters (α, β) (where $(1, 1)$ is the uniform distribution).

In all specifications, we pool the data across $N = 3$ and $N = 6$ bids, imposing the equality of the structural parameters across these two sets of bids. The reason for this is that, as argued above, without variation in the number of bidders, the Bayes Nash benchmark model where $\lambda \rightarrow \infty$ can explain any distribution of bids with a sufficiently flexible distribution of valuations. Hence, intuitively, to identify λ, one needs an extra source of variation in the data, such as that given by the variation in the number of bidders.

---

[22] Haile, Hortaçsu and Kosenok (2003) show that given a matrix of average payoffs in a game, there are infinitely many ways of choosing the idiosyncratic shock distribution such that any observed play probabilities can be rationalized within the QRE framework – *even if* one restricts the (zero mean) idiosyncratic shock terms to be (A) independently distributed, or (B) identically distributed across a player's actions. This implies that data on play probabilities do not contain any information about the average payoff matrix of a game if one is willing to be flexible about the idiosyncratic shock distribution within the classes (A) and (B).

In the estimation exercise we also have to discretize the strategy space of the bidders to calculate the QRE probabilities. We discretize all bids in the data set to their nearest dollar increment. We found that using 10 cent increments increases the computational requirements of the model considerably. Moreover, the likelihood functions (25) and (26) can not be computed analytically. Therefore, we use a simulated likelihood approach where we use 1,000 random draws from the latent distribution of valuations, $f(v|\omega)$ for each likelihood contribution.

In Table 7, we report the results of the three specifications. The risk neutral specification with uniformly distributed valuations reflects what we expect from the "overbidding" phenomenon, and yielded similar results to the risk neutral Bayes-Nash model by overestimating the upper support of the valuation distribution. The lower support of the distribution is also overestimated by the QRE specification. We believe this is driven by the fact that the expected payoffs of bidders with valuations in this region are close to zero, and thus dominated by the QRE noise term. Thus, the QRE assigns roughly equal probabilities to the different bidding strategies of bidders with low valuations – i.e. the econometric model can not reliably distinguish between valuations rationalizing low bids in the data.

Allowing for risk aversion (expectedly) improves the fit of the model. Notice that the upper support of the valuation distribution is estimated to be 27.75, closer to 30 than the risk neutral specification. This indicates that allowing for risk aversion, as before, can account for the overbidding phenomenon. Notice also that the estimated CRRA exponent, θ, is 0.27, which is very much in line with the estimates from the symmetric risk averse Bayes-Nash model. However, the lower support of the valuation distribution was once again overestimated as being 5.16.[23] With these support estimates, the Kolmogorov-Smirnov

---

[23] We note that the standard error estimates in this specification are somewhat suspect, since the Hessian of the log-likelihood function was ill-conditioned.

sup-distance between the "estimated" valuation distribution and the actual distribution was calculated to be 0.172. This point value is about twice as large as the K-S distances reported in Table 4, suggesting that the Bayes-Nash model provided a better estimate of the valuation distribution under this (sup-norm) metric.[24]

In the last specification, we relax the uniform distribution assumption by allowing for a beta distribution. The support estimates are very similar to those obtained from the uniform specification. The CRRA exponent, θ is estimated to be 0.25, which is once again very similar to the estimate obtained in the Bayes-Nash symmetric risk aversion model. However, the estimated beta parameters indicate a valuation distribution significantly different from the uniform. Moreover, the Kolmogorov-Smirnov distance of this beta distribution (evaluated at the point estimates of the parameters) from the actual distribution of valuations was 0.192, which was again larger than the K-S distances for the symmetric risk aversion estimates reported in Table 4.

We thus conclude that the QRE model with risk aversion provides similar results to the Bayes-Nash with symmetric risk aversion model; though it is unable to pin down the lower support of the valuation distribution correctly. Moreover, the estimator is more difficult to compute than the estimators of the previous sections since it requires solving a multivariate, nonlinear optimization problem.[25] The computational tractability of the QRE model requires a parsimonious, parametric specification. In field work, a priori knowledge of the distribution of private information may be a strong assumption.

7 Conclusion

In this paper, we attempt to provide some evidence on the usefulness of structural models of bidding in first-

---

[24] We should caution the reader that the comparison of the K-S distance across the Bayes-Nash and QRE models does not have a rigorous statistical interpretation.

[25] The likelihood iterations for the QRE model took several hours to converge within acceptable limits, whereas the nonparametric procedures used in previous sections were completed in seconds.

price auctions. If the researcher does not make parametric assumptions about the distribution of private information, his ability to reject or accept the theory using bid data alone is limited.

Table 8 summarizes our results regarding the distance between actual and estimated distribution of valuations yielded by the different estimation methods (we did not include the QRE results here since we do not have the results separately for the $N = 3$ and $N = 6$ cases, and the interpretation of the metrics is somewhat different due to the parametric nature of the estimation method used there). Our interpretation of these results is that when the number of bidders is sufficiently large, most of the methods do a reasonable job in uncovering the structural parameters. When the number of bidders is smaller, the results are more sensitive to the choice of method. The risk aversion model is closer to the actual distribution of valuations in almost all norms. In the L^1 and L^2 norms, the differences between the models are particularly large. The average distance between the estimated and actual valuations is \$1.39 for the risk aversion model compared to \$3.98 for the risk neutral model and \$5.09 for the learning model.

Just because the risk aversion model appears to perform the best when recovering bidder valuations from laboratory data does not imply that it should always be preferred when analyzing field data. The stakes in the laboratory experiments are fairly small and may not be indicative of behavior in the field. The identifying assumptions are more demanding for the risk aversion model since structure on the distribution of valuations is assumed constant across auctions. In some empirical applications, the cost of these assumptions may outweigh the benefits.

Despite the fact that, in our chosen metrics, 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 risk aversion produces better estimates than models in which bidders are risk neutral.

8 Appendix: Derivation of the Limiting Distribution of the Modified KS statistic

We will now show that the modified KS-statistic proposed in section 3.2, $M\tilde{K}S_T$, has a nondegenerate limiting distribution. Recall that the statistic is given by:

$$M\tilde{K}S_T = \sqrt{T} \sup_{v \in [\underline{v}, \bar{v}]} |\frac{1}{T}\sum_{t=1}^{T} \Lambda(\hat{v}_t - v) - F(v)|$$

Let

$$\begin{aligned} L(v) &= \frac{1}{T}\sum_{t=1}^{T} \Lambda(\hat{v}_t - v) - F(v) \\ &= \underbrace{\left[\frac{1}{T}\sum_{t=1}^{T} \Lambda(v_t - v) - F(v)\right]}_{L_1(v)} + \underbrace{\left[\frac{1}{T}\sum_{t=1}^{T} \Lambda(\hat{v}_t - v) - \frac{1}{T}\sum_{t=1}^{T} \Lambda(v_t - v)\right]}_{L_2(v)} \end{aligned}$$

The first term in brackets, $L_1(v)$, is $O_p(\frac{1}{\sqrt{T}})$ uniformly in v (the empirical cdf converges at rate $\frac{1}{\sqrt{T}}$). Hence $\sqrt{T}L_1(v)$ has the standard limiting Gaussian bridge distribution of the empirical cdf.

We now look at the convergence properties of $\sqrt{T}L_2(v)$. To do this, take a term-by-term Taylor approximation around v_t:

$$\begin{aligned}
\sqrt{T}L_2(v) &= \sqrt{T}\frac{1}{T}\sum_{t=1}^{T}\lambda(v_t - v)(\hat{v}_t - v_t) + o_p(1) \\
&= \sqrt{T}\frac{1}{T}\sum_{t=1}^{T}\lambda(v_t - v)\frac{1}{N-1}\left(\frac{\hat{G}(b_t)}{\hat{g}(b_t)} - \frac{G(b_t)}{g(b_t)}\right) + o_p(1) \\
&= \sqrt{T}\frac{1}{T}\sum_{t=1}^{T}\lambda(v_t - v)\frac{1}{N-1}\left(\frac{\hat{G}(b_t)}{g(b_t)} - \frac{G(b_t)}{g(b_t)} - \frac{\hat{G}(b_t)}{g^2(b_t)}\left[\hat{g}(b_t) - g(b_t)\right]\right) + o_p(1) \\
&= \sqrt{T}\frac{1}{T}\frac{1}{N-1}\sum_{t=1}^{T}\lambda(v_t - v)\left\{\begin{array}{c}
\underbrace{\frac{\hat{G}(b_t) - G(b_t)}{g(b_t)}}_{T_1} - \underbrace{\frac{\hat{G}(b_t) - G(b_t)}{g^2(b_t)}\left[\hat{g}(b_t) - g(b_t)\right]}_{T_2} \\
\underbrace{-\frac{G(b_t)}{g^2(b_t)}\left[\hat{g}(b_t) - g(b_t)\right]}_{T_3}
\end{array}\right\} + o_p(1)
\end{aligned}$$

where the first line follows from the Taylor approximation, second line from a substitution of the formula for $\hat{v}_t$, the third term from a Taylor approximation around $g(b_t)$. The fourth term is a rearrangement of the third.

Now let's look at the individual terms T_1, T_2 and T_3. As above, $\hat{G}(b_t)$ converges to $G(b_t)$ at rate $O_p(\frac{1}{\sqrt{T}})$. Hence the contribution of the T_1 term to the sum is $o_p(1)$. Similarly, the contribution of T_2 terms is also $o_p(1)$. Thus we can write:

$$\begin{aligned}
\sqrt{T}L_2(v) &= \sqrt{T}\frac{1}{T}\sum_{t=1}^{T}\lambda(v_t - v)\frac{G(b_t)}{g^2(b_t)}\left[g(b_t) - \hat{g}(b_t)\right] + o_p(1) \\
&= \frac{1}{\sqrt{T}}\frac{1}{N-1}\sum_{t=1}^{T}\lambda(v_t - v)\frac{G(b_t)}{g^2(b_t)}\left[E\hat{g}(b_t) - \hat{g}(b_t)\right] + o_p(1) \\
&= EQ(v) - Q(v) + o_p(1)
\end{aligned}$$

where we used the bias formula for the kernel density estimator. Now substitute in for $\hat{g}(b_t) =$

$\frac{1}{T}\sum_{s=1}^{T}\frac{1}{h_T}K(\frac{b_t-b_s}{h_T})$, and noting that $v_t = v(b_t)$, we get a U-statistic representation for the statistic $Q(v)$ at each value of v:

$$Q(v) = \sqrt{T}\frac{1}{T^2}\frac{1}{N-1}\sum_{t=1}^{T}\sum_{s=1}^{T}\lambda(v(b_t)-v)\frac{G(b_t)}{g^2(b_t)}\frac{1}{h_T}K(\frac{b_t-b_s}{h_T})$$

To establish the asymptotic normality of this statistic for each v, we use the Hajek projection (van der Vaart (1998), p. 162) to project $Q(v)$ onto a one dimensional statistic. To do this, define:

$$P(v) = \sqrt{T}\frac{1}{T}\frac{1}{N-1}\int_{b_t}\sum_{s=1}^{T}\lambda(v(b_t)-v)\frac{G(b_t)}{g^2(b_t)}K(\frac{b_t-b_s}{h_T})g(b_t)\frac{1}{h_T}db_t$$

which is a one dimensional statistic. By the projection formula, $EQ(v)-Q(v) = EP(v)-P(v)+o_p(1)$. Then, for $b_t = b_s + h_T\varepsilon$:

$$\begin{aligned} P(v) &= \sqrt{T}\frac{1}{T}\frac{1}{N-1}\sum_{s=1}^{T}\int_{\varepsilon}\lambda(v(b_s+h_T\varepsilon)-v)\frac{G(b_s+h_T\varepsilon)}{g(b_s+h_T\varepsilon)}K(\varepsilon)d\varepsilon \\ &= \frac{1}{\sqrt{T}}\frac{1}{N-1}\sum_{s=1}^{T}\lambda(v(b_s)-v)\frac{G(b_s)}{g(b_s)}\int_{\varepsilon}K(\varepsilon)d\varepsilon + o_p(1) \\ &= \frac{1}{\sqrt{T}}\frac{1}{N-1}\sum_{s=1}^{T}\lambda(v(b_s)-v)\frac{G(b_s)}{g(b_s)} + o_p(1) \end{aligned}$$

where the second line follows from Pagan and Ullah (1999), Lemma (A.55). For each v, this is an asymptotically normal statistic with variance (b_s are independent):

$$\begin{aligned} Var(P(v)) &= \frac{1}{N-1}E\left[\left(\lambda(v(b_s)-v)\frac{G(b_s)}{g(b_s)}\right)^2\right] \\ &= \frac{1}{N-1}\int\left(\lambda(v(b_s)-v)\frac{G(b_s)}{g(b_s)}\right)^2 g(b_s)db_s \end{aligned}$$

which, as can be seen, does not depend on T (or h_T).

The covariance term $Cov(P(v), P(v'))$ is given by:

$$Cov(P(v), P(v')) = Cov\left(\frac{1}{\sqrt{T}}\frac{1}{N-1}\sum_{s=1}^{T}\lambda(v(b_s)-v)\frac{G(b_s)}{g(b_s)}, \frac{1}{\sqrt{T}}\frac{1}{N-1}\sum_{s=1}^{T}\lambda(v(b_s)-v')\frac{G(b_s)}{g(b_s)}\right)$$

which once again won't depend on T.

Hence, the results obtained above show that, under the null hypothesis:

$$\sqrt{T}\sup_{v\in[\underline{v},\bar{v}]}|\frac{1}{T}\sum_{t=1}^{T}\Lambda(\hat{v}_t - v) - F(v)| \xrightarrow{d} \sup_{v\in[\underline{v},\bar{v}]}\{G(v)+H(v)\}$$

where $G(v)$ is a zero-mean Gaussian process on $v \in [\underline{v}, \bar{v}]$ with:

$$Var(G(v)) = \lim_{T\to\infty} Var(P(v))$$

and

$$Cov(G(v), G(v')) = \lim_{T\to\infty} Cov(P(v), P(v'))$$

and $H(v)$, which is the limiting distribution of $\sqrt{T}L_1(v)$, is also a Gaussian process very similar to the limiting distribution of the empirical cdf (and indeed is the empirical cdf when $h' \to 0$).

Hence the limiting distribution of the modified KS-statistic is the distribution of the supremum of a Gaussian process defined over $v \in [\underline{v}, \bar{v}]$, a well-defined object.

9 References

[1] Bajari, P. and Hortaçsu, A. (2003) "The Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions." *RAND Journal of Economics*, 34, pp. 329-355.

[2] Bajari, P. and Ye, L. (2003) "Deciding Between Competition and Collusion", *Review of Economics and Statistics,* 85(4), pp. 971-989.

[3] Bray, M. (1982) "Learning, Estimation, and the Stability of Rational Expectations." *Journal of Economic Theory*, 26(2), pp. 318-339.

[4] Campo, S., Guerre, E., Perrigne, I. and Vuong, Q. (2000) "Semiparametric Estimation of First-Price Auctions with Risk Averse Bidders", University of Southern California Working Paper.

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

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

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

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

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

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

[11] Donald, S. and Paarsch, H. J. (1993) "Piecewise Pseudo-maximum Likelihood Estimation in Empirical Models of Auctions", *International Economic Review,* 34(1), pp. 121-48.

[12] Donald, S. and Paarsch, H.J. (1996) "Identification, Estimation, and Testing in Parametric Empirical Models of Auctions within the Independent Private Values Paradigm", *Econometric Theory,* 12(3), pp. 517-67.

[13] Dyer, D., Kagel, J. H., Levin, D. (1989) "Resolving Uncertainty About Numbers of Bidders in Independent Private Value Auctions: An Experimental Analysis," *The RAND Journal of Economics,* 20(2), pp. 268-279.

[14] Elyakime, B., Laffont, J.J., Loisel, O., and Vuong, Q. (1994) "First-Price Sealed-Bid Auctions with Secret Reservation Prices," *Annales d'Economie et de Statistique,* 0(34), pp. 115-41.

[15] Flambard, V. and Perrigne, I. (2002) "Asymmetry in Procurement Auctions: Some Evidence from Snow Removal Contracts", University of Southern California Working Paper.

[16] Genesove, D and Mullin, W. P. (1998), "Testing Static Oligopoly Models: Conduct and Cost in the Sugar Industry, 1890-1914," *RAND Journal of Economics*, 29(2), pp. 355-377.

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

[18] Goeree, J., Holt, C. and Palfrey, T. (2002) "Quantal Response Equilibrium and Overbidding in Private-Value Auctions," *Journal of Economic Theory*, 104(1), pp. 247-272.

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

[20] Haile, P., Hong, H., Shum, M. (2003) "Nonparametric Tests for Common Values In First-Price Sealed Bid Auctions," Yale University Working Paper.

[21] Haile, P., Hortaçsu, A., Kosenok, G. (2003) "On the Empirical Content of Quantal Response Equilibrium," Yale University Working Paper.

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

[23] Harrison, G. W. (1990) "Risk Attitudes in First-Price Auction Experiments: A Bayesian Analysis," *Review of Economics and Statistics,* 72, pp. 541-46.

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

[25] Hendricks, K., Pinkse, J. and Porter, R. (2003) "Empirical Implications of Equilibrium Bidding in First-Price Symmetric, Common Value Auctions", *Review of Economic Studies,* 70, pp. 115-146.

[26] Holt, C. A. (1980) "Competitive Bidding for Contracts Under Alternative Auction Procedures," *Journal of Political Economy*, 88, pp. 433-445.

[27] Horowitz, J. (2001) "The Bootstrap," *Handbook of Econometrics, Vol.5*, J.J. Heckman and E.E. Leamer, eds., Elsevier Science B.V., 2001, Ch. 52, pp. 3159-3228.

[28] Hortaçsu, A. and Puller, S. (2004) "Testing Strategic Models of Firm Behavior in Restructured Electricity Markets: A Case Study of ERCOT", University of Chicago Working Paper.

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

[30] Kagel, J.H. (1995) "Auctions: A Summary of the Experimental Research." In *The Handbook of Experimental Economics,* Kagel, J. H. and Roth, A. E. (eds.). Princeton University Press, Princeton, pp. 501-535.

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

[32] Laffont, J.J. and Vuong, Q. (1996) "Structural Analysis of Auction Data", *American Economic Review,* 86(2), 414-20.

[33] Marcet, A. and Sargent, T.J. (1989) "Convergence of Least Squares Learning Mechanisms in Self Referential Linear Stochastic Models," *Journal of Economic Theory*, 48, pp. 337-368.

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

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

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

[37] Ockenfels, A., Roth, A. (2003), "Late and Multiple Bidding in Second Price Internet Auctions: Theory and Evidence Concerning Different Rules for Ending an Auction," *Games and Economic Behavior* (forthcoming).

[38] Paarsch, H. J. (1992) "Deciding Between the Common and Private Value Paradigms in Empirical Models of Auctions," *Journal of Econometrics*, 51 (1-2), 191-215.

[39] Pagan, A., Ullah, A. (1999), *Nonparametric Econometrics*, Cambridge University Press.

[40] Pesendorfer, M. and Jofre-Bonet, M. (2003) "Estimation of a Dynamic Auction Game", *Econometrica*, 71(5), pp. 1443-1489.

[41] Plott, C. and Salmon, T. (2004) "The Simultaneous, Ascending Auction: Dynamics of Price Adjustment in Experiments and in the U.K. 3G Auction," *Journal of Economic Behavior and Organization*,

53(3), pp. 353-383.

[42] Politis, D., Romano, J. and Wolf, M. (1989) *Subsampling*. Springer Series in Statistics.

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

[44] van der Vaart, A. (1998) *Asymptotic Statistics*. Cambridge University Press, Cambridge.

[45] Wolfram, C.D. (1999), “Measuring Duopoly Power in the British Electricity Spot Market,” *American Economic Review*, 89(4), pp. 805-826.

Table 1: Comparison of Observed versus Equilibrium Bidding Behavior

VARIABLE	#OBS	MEAN	STD	25 %	50 %	75 %
Observed-Nash Bid, *N=3*	414	2.44	1.86	1.05	2.46	3.84
Observed-Nash Bid, *N=6*	414	0.65	1.05	0.03	0.56	1.42
Expected Profit, *N=3*	414	1.33	1.61	0.08	0.62	2.05
Expected Profit, *N=6*	414	0.58	1.05	0.01	0.05	0.57
Optimization Error, *N=3*	414	0.40	0.49	0.03	0.19	0.63
Optimization Error, *N=6*	414	0.12	0.26	0.00	0.01	0.12

Table 2: Estimation Results for the Risk Neutral Model (p-values in parentheses)

	N=3	*N=6*
K-S statistic (25-75% of valuation support)	0.1691 (0.000)	0.0700 (0.1480)
Reject equality of distributions? (@ 5%)	Yes	No
K-S statistic (5-95% of valuation support)	0.1925 (0.000)	0.0672 (0.0720)
Reject equality of distributions? (@ 5%)	Yes	No
L^1 Norm	3.981	1.067
L^2 Norm	5.317	1.554

Table 3: Risk Aversion Parameter Estimates

	θ	*OLS 5-95% CONFIDENCE INTERVAL*
Entire sample	0.1580	[0.1382,0.1779]
5-95% of valuation support	0.2237	[0.2074,0.2399]
25-75% of valuation support	0.2474	[0.2303,0.2646]

Table 4: Estimation Results for the Symmetric Risk Aversion Model (p-values in parentheses)

	N=3	*N=6*	*JOINT TEST*
K-S statistic (25-75% of valuation support)	0.0342	0.0361	0.0703
	(0.56)	(0.48)	(0.52)
Reject equality of distributions? (@ 5%)	No	No	No
K-S statistic (5-95% of valuation support)	0.0618	0.0789	0.1407
	(0.42)	(0.094)	(0.2640)
Reject equality of distributions? (@ 5%)	No	No	No
K-S statistic (entire sample)	0.0867	0.0893	0.1760
	(0.14)	(0.014)	(0.0580)
Reject equality of distributions? (@ 5%)	No	Yes	No
L^1 Norm	1.387	1.344	
L^2 Norm	2.084	1.862	

Table 5: Percentiles for Actual and Estimated Valuations, Symmetric Risk Aversion Model

PERCENTILE	ACTUAL VALUATIONS	ESTIMATED VALUATIONS, *N=3*	ESTIMATED VALUATIONS, *N=6*
10	3.94 (2.99,4.98)	3.39 (2.45,4.36)	3.32 (2.22,4.21)
20	7.08 (6.00,7.81)	6.19 (5.31,7.33)	6.36 (5.22,7.33)
30	9.94 (8.55,11.20)	9.04 (7.86,10.15)	9.07 (7.63,10.49)
40	12.95 (11.36,14.10)	12.11 (10.76,13.44)	11.64 (10.57,13.38)
50	15.81 (14.16,17.26)	14.68 (13.45,16.26)	14.87 (13.57,16.35)
60	19.04 (17.36,20.11)	17.54 (16.36,19.26)	17.62 (16.49,18.95)
70	22.01 (20.27,23.21)	20.87 (19.35,21.94)	20.82 (19.05,21.70)
80	25.09 (23.76,25.93)	22.90 (22.15,23.50)	23.08 (22.28,23.90)
90	27.69 (26.98,28.37)	25.28 (24.54,26.00)	24.59 (24.10,25.04)
99	29.99	30.11	27.62

Notes: This table provides the estimated percentiles for the actual distribution of valuations and the estimated distribution of valuations from the risk aversion model. Ninety-five percent confidence intervals for the percentiles are listed in the parentheses.

Table 6: Estimation Results for the Learning Model

	N=3	*N=6*
K-S distance	0.2500	0.0784
L^1 Norm	5.091	1.466
L^2 Norm	6.745	2.012

Table 7: Estimation Results for the QRE Model

	RISK NEUTRAL (UNIFORM DIST)	*RISK AVERSION (UNIFORM DIST)*	*RISK AVERSION (BETA DIST)*
Vlower	3.48 (0.08)	5.16 (0.001)	4.39 (0.08)
Vupper	36.16 (0.19)	27.75 (1.05)	27.66 (0.19)
λ	15.68 (0.16)	17.36 (1.00)	11.32 (0.10)
θ	-	0.27 (1.29)	0.25 (0.03)
α	-	-	1.44 (0.004)
β	-	-	1.10 (0.01)
Log-likelihood	-2873.3	-2810.4	-2806.3

Notes: Data was pooled across N=3 and N=6 bids, imposing the equality of parameter values underlying these bids. All estimates were obtained by discretizing the bids to \$1 increments. Numerical integration to compute the likelihood of each observation was performed using 1,000 Monte-Carlo draws from the underlying valuation distribution. Standard errors, reported in parantheses, were obtained using the numerically computed Hessian. The Hessian was badly scaled in the “risk aversion with uniform distribution” case.

Table 8: Comparison of Models

	RISK NEUTRAL	*RISK AVERSE*	*LEARNING*
N=3			
K-S Statistic (on 5-95% of valuation support)	0.1925	0.0618	0.2500
Reject equality of distributions? (@ 5%)	Yes	No	Yes
L^1 Norm	3.981	1.387	5.091
L^2 Norm	5.317	2.084	6.745
N=6			
K-S Statistic (on 5-95% of valuation support)	0.0672	0.0789	0.0784
Reject equality of distributions? (@ 5%)	No	No	No
L^1 Norm	1.067	1.34	1.466
L^2 Norm	1.554	1.862	2.012

Figure 1: Comparison of Equilibrium and Actual Bids.

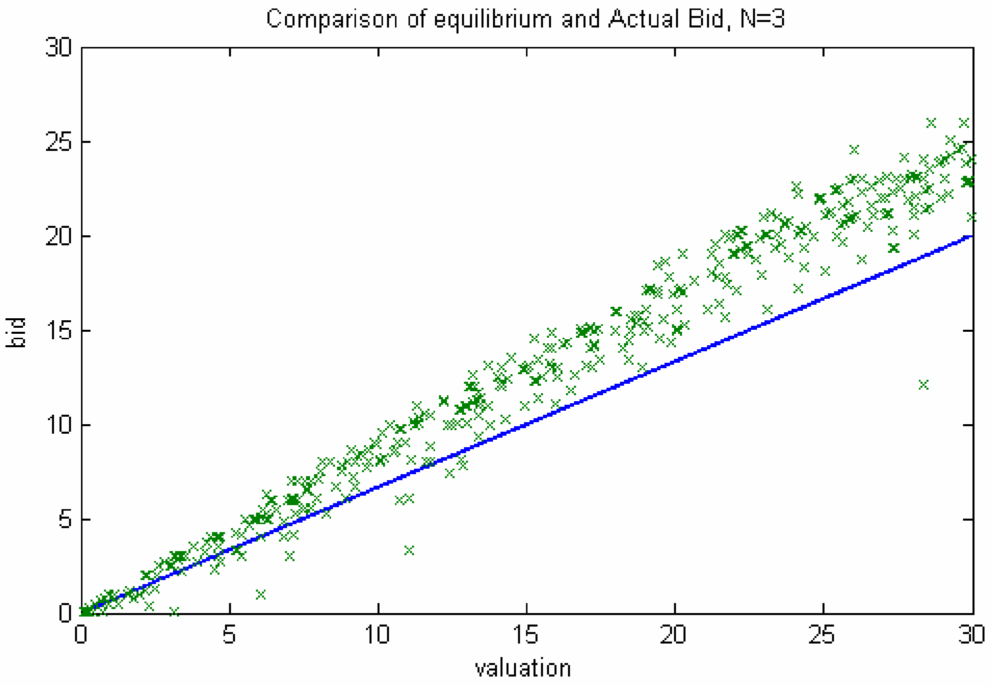

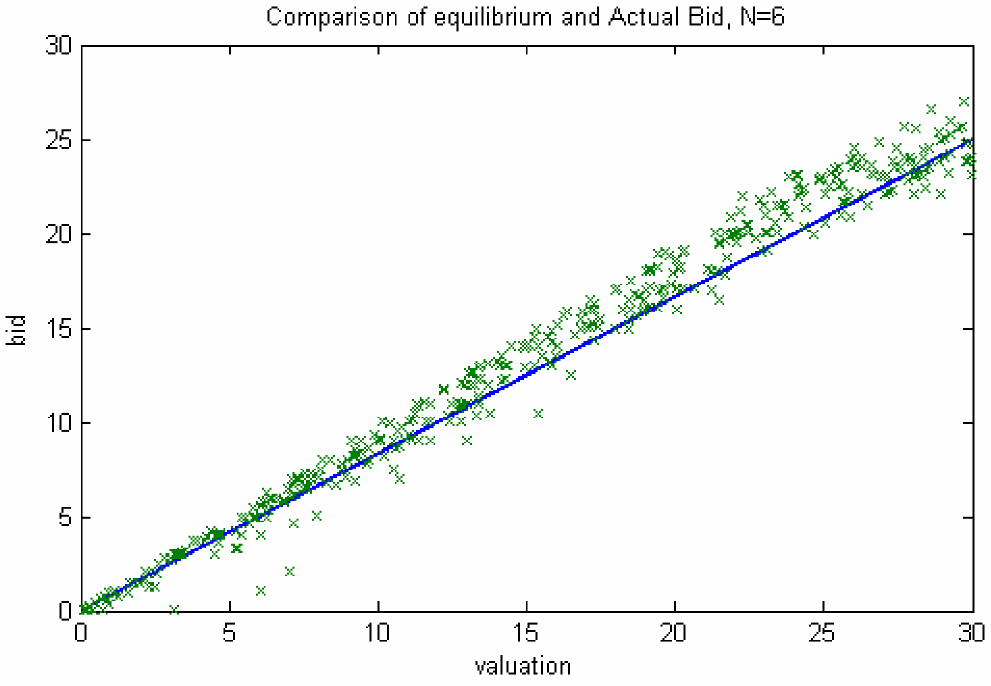

Figure 2: Histograms of Estimated and Actual Valuations, Risk Neutral Model.

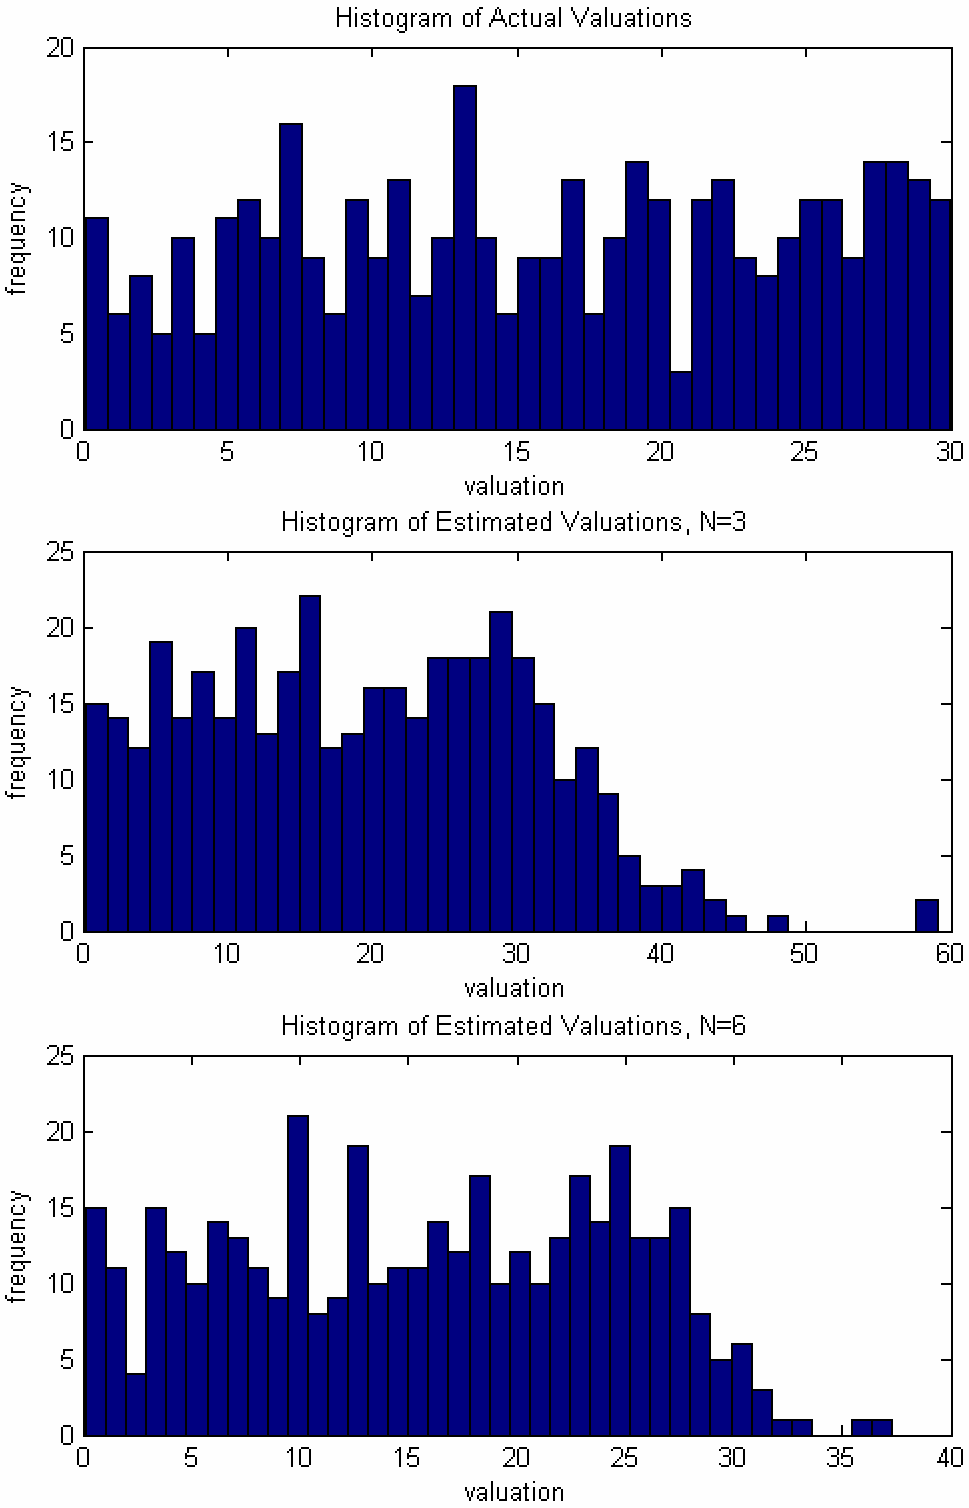

Figure 3: Histograms of Estimated and Actual Valuations, Risk Aversion Model.

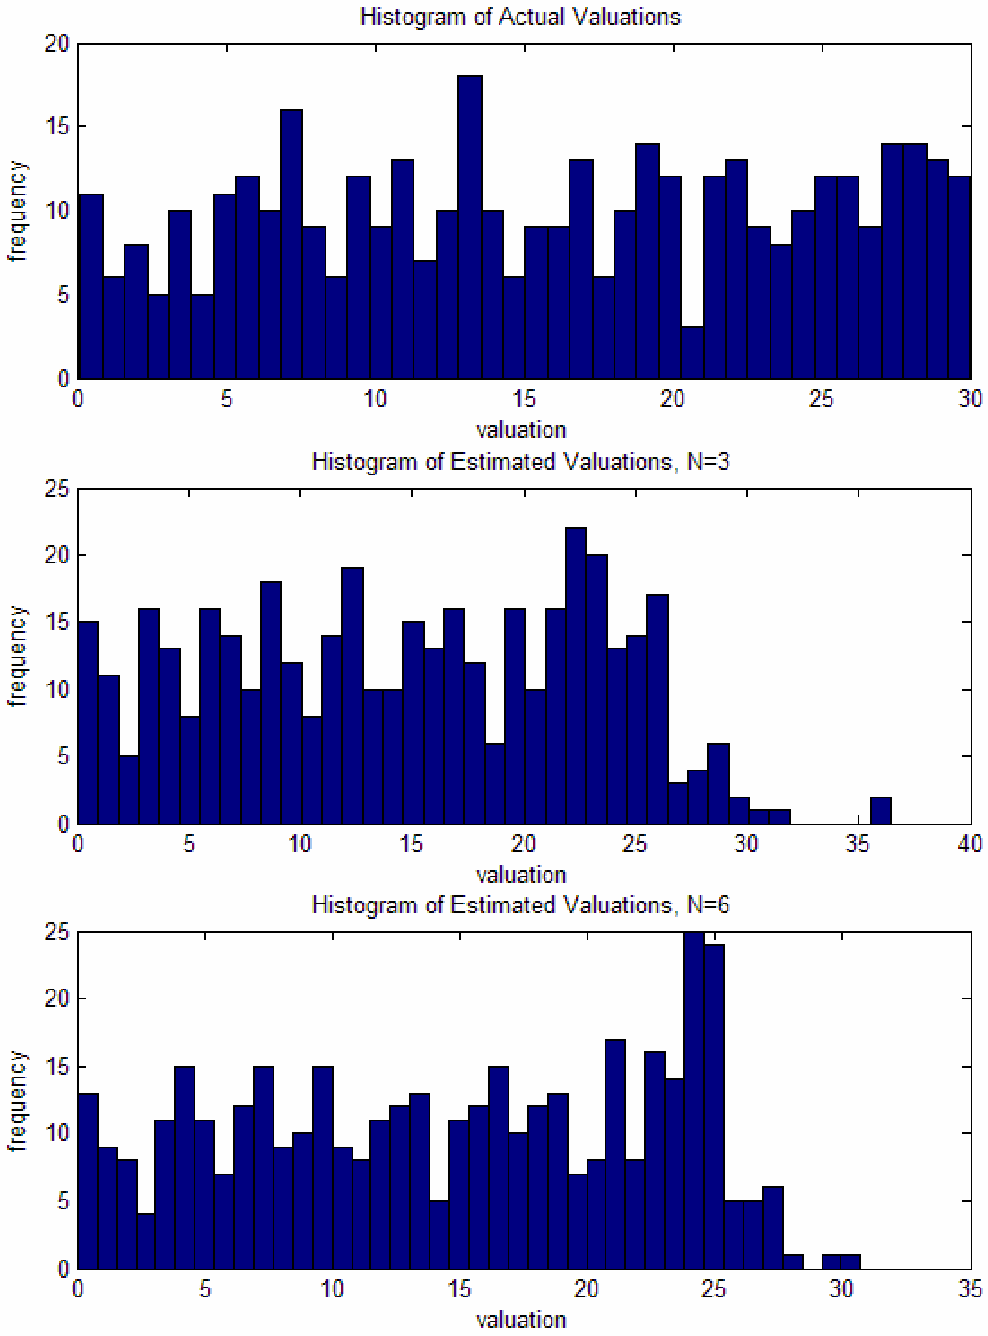

Histogram of Actual Valuations

Histogram of Estimated Valuations, N=3

Histogram of Estimated Valuations, N=6

Figure 4: Histograms of Estimated and Actual Valuations, Learning Model.