Structural inferences from first-price auction data*

Paul Pezanis-Christou

Bureau d'Economie Théorique et Appliquée
Université Louis Pasteur
61, Avenue de la Forêt Noire
67085 Strasbourg, France

Andrés Romeu Santana

Departamento de Fundamentos del Análisis Económico,
Universidad de Murcia,
Campus Espinardo
30100 Murcia, España

January 2007

PRELIMINARY VERSION

Abstract

We use structural econometric methods to assess the symmetric Bayes-Nash equilibrium model of bidding with data from first-price auction experiments. We analyze the data from both empirical and experimental perspectives. Our study focuses on the sensitivity of structural inferences to some basic behavioral assumptions and to the quality of the information available to the field researcher. It shows that once behavior is diagnosed "out-of-equilibrium", structural estimates become highly sensitive to the information available, and do not improve with the quality of information. We also devise econometric procedures to test the symmetric Bayes-Nash model for homogenous or heterogeneous constant relative risk averse bidders in experimental contexts, and report an overwhelming rejection of this model to explain bidding behavior at first-price auctions with independent private values.

Key Words: first-price sealed bid auctions, private independent values, experiments, structural econometrics, constant relative risk aversion, optimal reserve price.

JEL Classification Numbers: C9, D44

---

* We thank Olivier Armantier, Jacob Goeree, François Laisney, Harry Paarsch and seminar participants at GREQAM, the Institute for Economic Analysis (CSIC, Barcelona), the Max Planck Institute (Jena) and the Economic Science Association conferences in Strasbourg and Tucson for helpful comments and discussions; and James Cox, John Kagel and James Walker for access to their data. Financial support from the European Commission through a EU-TMR ENDEAR Network grant (FMRX-CT98-0238), the *Ministerio de Ciencia y Tecnologia* (Grant BEC-2001-0980) is gratefully acknowledged. Remaining errors are ours.

	Pezanis-Christou	Romeu Santana
e-mail	ppc@cournot.u-strasbg.fr	aromeu@um.es
Telephone	+33.390.242.075	+34.968.367.909

1. Introduction

The growing use of auctions to allocate assets has been paralleled in recent years by the development of an impressive literature on the structural estimation of auction models (see Paarsch and Hong, 2006, for an up-to-date review of the literature). In its essence, the structural approach to data is based on the assumption that the variables of interest are in equilibrium, as is the case when studying macroeconomic equilibrium models of the economy. In the context of games with incomplete information such as auctions, however, equilibrium models can difficult to assess because the relevant data are usually not observable (e.g., the number of active bidders, the functional form of bidders' distribution(s) of private values, their risk preferences and the realizations of their private values). Therefore, rather than assessing equilibrium models, the structural approach to auction data typically consists in recovering the unobserved information by assuming bidders to act in a given equilibrium. Inference based on this approach --- henceforth, structural inference --- is fundamental for policy recommendations such as setting a revenue enhancing optimal reserve price, or implementing an alternative selling mechanism. But the success of such recommendations, which may imply non-negligible money transfers from one side of a market to another, remains highly conditional on the assumption that the recovered information is correct or, equivalently, that the observed bids or prices are in the presumed equilibrium.

The first goal of this paper is to assess the sensitivity of structural inferences to some basic assumptions of the model used to draw these inferences and to the information that is typically available to the field researcher. To this end, we use laboratory data and a methodology inspired from Laffont, Ossard and Vuong (1995) for the analysis of symmetric first-price auctions with private independent values. We motivate the use of experimental rather than field data by the fact that in experiments, the auction fundamentals are perfectly known to the researcher and are made explicit to the participating bidders. Our choice to study symmetric auctions further keeps the analysis simple as it requires that the bidders' private values are all drawn (with replacement) from the same probability distribution. This assumption is crucial to the determination of equilibrium predictions for the first-price auction format, and offers some simple behavioral hypotheses to be tested for the observed behavior to be a Symmetric Bayes-Nash Equilibrium (SBNE) one. Thus, by controlling the exhaustiveness of the analysis (i.e., prior check of a symmetric / homogenous bidding behavior) and the quality of the information available (e.g., information on all bids or only on winning bids), we are able to assess the SBNE model and to check how inferences and predictions are affected by various assumptions that underly this model.

So far we assumed an empirical setup in that we discarded information about the bidders' value realizations from analysis. The paper's second goal is to build a structural model that takes explicit account of this information for the analysis of experimental data. In first-price auction experiments, behavior is typically characterized by a significant overbidding (i.e., bidding above the risk neutral Nash equilibrium prediction) and heterogeneity; both patterns are usually found to be well or best explained by a SBNE model that assumes heterogeneous constant relative risk attitudes (Cox *et al.*, 1982, 1983, 1988, 1996). These conclusions, however, are based on the estimation of reduced-form expressions of equilibrium behavior, and have raised some controversy in the literature because they fail to explain patterns such as significantly negative intercept terms in the estimation of individual bid functions and less-than proportional bidding at higher values in symmetric settings (Kagel, 1995, Chen and Plott, 1998).[1] It appears to us that part of this controversy results from the lack of a structural assessment of the SBNE model in experimental conditions. We therefore identify conditions that allow a straightforward assessment of this model and of its augmented version for heterogeneous risk attitudes with experimental data.

Our analysis complements the one of Bajari and Hortaçsu (2005) who estimate structural econometric models with experimental first-price auction data.[2] Their study concentrates on fitting various behavioral models in the aggregate and on comparing the non-parametrically estimated values to the actual ones. They show in particular that the SBNE model for homogenous constant relative risk averse bidders organizes the data almost better than models that assume either risk neutral SBNE bidders (with and without adaptive learning) or a QRE model with constant relative risk aversion. Our analysis deals instead with assessing (i.e., testing and evaluating) the SBNE model under various information conditions that pertain to either empirical or experimental setups. It reveals in particular that when behavior is diagnosed as

[1] A vivid debate about overbidding at first-price auctions and the risk aversion hypothesis was also triggered by Harrison's (1989) *flat maximum* critique --- cf. Cox, Smith and Walker (1992), Friedman (1992), Harrison (1992), Kagel and Roth (1992) and Merlo and Schotter (1992). Numerous other studies have been conducted to bracket the effect of risk aversion on bidding behavior within a Bayes-Nash equilibrium context, among which Harrison (1990), Goeree, Holt and Palfrey (2002) and Armantier and Treich (2005). Ockenfels and Selten (2005) and Neugebaeur and Selten (2006) study bidding dynamics and the effect of information feedback at first-price auctions. They show that the observed behavior is not compatible with the (risk-averse) Bayes-Nash prediction which supposes a stationary behavior.

[2] See also Armantier (2002) who uses non-parametric structural methods to analyze experimental data on affiliated value auctions and decide which of the *common value* or the *private independent value* framework applies best.

being “out-of-equilibrium”, structural inferences become highly sensitive to the information available. A comparison of our (parametric) estimates to those of Bajari and Hortaçsu further reveals important differences that put in perspective the use of non-parametric structural methods for the analysis of experimental auction data.

In the next section we outline the econometric procedures used to test the SBNE model in empirical and experimental conditions. Section 3 reports and discusses the outcomes and Section 4 concludes.

2. Theory and Model Specification

2.1. *Symmetric Bayes-Nash Equilibrium Behavior*

In a first-price sealed-bid auction, n bidders compete for the purchase of a single commodity which is awarded to the highest bidder for a price equal to her/his bid. Each bidder is assumed to receive a private reservation value v_i, which is an independent draw from a distribution F with support $[\underline{v};\overline{v}]$ and density f. Bidders have the same utility function $u(\cdot)$ on monetary payments, with $u' > 0$, $u'' \leq 0$ and $u(0) = 0$. A bidder with value v_i who submits a bid b_i has a utility of winning the auction equal to $u(v_i - b_i)$. The number of bidders, F and $u(\cdot)$ are common knowledge but the value realizations v_i are private information. In this context, a bidding strategy $b(\cdot)$ is a SBNE strategy if for all valuations, it is a best response for bidder i to use $b(\cdot)$ if all bidders $j \neq i$ also use $b(\cdot)$. Maskin and Riley (2000) further show that if $b(\cdot)$ is a bidder's best response then it is monotone increasing in valuations. Therefore, if $b^{-1}(b_i)$ stands for the inverse of $b(\cdot)$ then bidder i's expected payoff is defined as

$$U(v_i,b_i) = u(v_i - b_i)[F(b^{-1}(b_i))]^{n-1}. \qquad (1)$$

Using the Bayes-Nash best-response first-order condition and imposing a symmetric behavior (i.e., $b_i = b(v_i), \forall i = 1,...,n$) yields the following nonlinear first order differential equation:

$$b'(v_i) = \frac{d\ln[F(v_i)]^{n-1}}{d\ln u(v_i - b(v_i))} \qquad \text{with } b(0) = 0 \ . \qquad (2)$$

This equation indicates that in equilibrium the slope of the bid function is equal to the elasticity of substitution between the probability of winning and bidder *i*'s payoff. As risk neutrality is characterized by $u(w) = w$, overbidding could be explained in terms of risk aversion since by concavity $d \ln u(w) < d \ln w, \forall w > 0$; i.e., the slope of the bid function is steeper with risk aversion than with risk neutrality. Following Holt (1980) and Cox, Smith and Walker (1982), we assume bidders to display homogenous constant relative risk averse preferences so that $u(w) = w^r$ for $i = 1,...,n$ and where $1 - r \neq 0$ represents the buyers' common Arrow-Pratt index of constant relative risk aversion.[3] With such preferences, equation (2) yields the following SBNE bid function

$$b_i = \varphi(v_i; F, \underline{v}, r) \equiv v_i - \int_{\underline{v}}^{v_i} \left[\frac{F(s)}{F(v_i)} \right]^{\frac{n-1}{r}} ds\,, \tag{3}$$

which is linear in values if the latter are uniformly distributed in $[\underline{v}; \overline{v}]$. The seller's expected revenue being equal to the expected winning bid, it is determined as

$$ER(p^*; F, n, r) = \int_{p^*}^{\overline{v}} \varphi(F, s, p^*, r)\; dF(s)^n\,, \tag{4}$$

where the integration is taken with respect to the distribution of the sample's maximum valuation, $F(s)^n$, and where $p^* \geq \underline{v}$ stands for the seller's reserve price provided that her reservation value for the commodity is lower than $\underline{v}$. Clearly, p^* is optimal if it maximizes equation (4). Actually, p^* is the solution to $0 = p^* - \left[1 - F(p^*)\right] / f(p^*)$ when $r = 1$, and it must be numerically determined when $r \neq 1$. An interesting property of this optimal reserve price is that it decreases with risk aversion (Riley and Samuelson, 1981, Proposition 5). The intuition for this is that highly risk averse bidders submit bids that are almost equal to their private values so that the rents the seller extracts from bidders by setting an optimal reserve price do not outweigh the risk of lowering competition. Hence it is of crucial importance for the seller

[3] Such preferences suit better the study of risk aversion in experimental settings than constant absolute risk averse (CARA) preferences because they encompass risk neutrality as a special case.

to know the type of preferences buyers have, for setting an optimal reserve price that assumes risk neutral bidders is not optimal if bidders are risk averse.

2.2 *Empirical and Experimental Frameworks*

We will consider two estimation frameworks, both of which assume the actual number of participating bidders to be common knowledge for the bidders and to be perfectly known by the researcher. The first framework pertains to empirical studies in that it assumes the researcher to be unaware of the bidders' private values and to have only partial information about the functional form of the distribution F. The second framework pertains to experimental studies in that the researcher is assumed to know both the functional form of F and the realizations of bidders' private values. Henceforth, we will refer to these frameworks as the empirical and the experimental framework, respectively. It is assumed that data have been collected from a series of T first-price auctions (indexed by $t = 1,\ldots,T$) of a homogenous commodity. Typically, in an empirical framework, *t* would represent different days or sampling sessions, while in an experimental framework *t* would be repetitions or rounds of the auction experiment.

We already note that although each framework implies its own econometric model specification, both will assume that the distribution of valuations *F* to belong to the Beta family of distribution functions indexed by only two parameters (θ_1 and θ_2) which nests a wide variety of distributions, suchas the uniform as well as other asymmetric distributions with only two parameters.[4] Since the domain of the Beta distribution is the unit interval, we normalize the support of values to this domain such that $\underline{v} = 0$ and $\overline{v} = 1$. To further lighten the notation, we will suppress the $\underline{v}$ argument of $\varphi(v_i; F, \underline{v}, r)$ and, since the Beta distribution is indexed by parameter vector $\theta = (\theta_1, \theta_2)$, we will refer to it as F_θ so that the bid function takes the expression $\varphi(v_i; \theta, r)$.

Our motivation for using a parametric rather than a non-parametric approach *à la* Guerre, Perrigne and Vuong (2000) or Campo *et al.* (2000) is twofold. First, as we use experimental auction data to tackle issues concerning the structural analysis of first-price auctions data, knowing the bidders' true distribution of values is important as it provides a simple way to check how the knowledge of this distribution and of its realizations affects ones structural inferences.

---

[4] The functional form of Beta distribution is: $F(v; \theta_1, \theta_2) = \frac{\Gamma(\theta_1 + \theta_2)}{\Gamma(\theta_1)\Gamma(\theta_2)} \int_0^v t^{\theta_1 - 1}(1-t)^{\theta_2 - 1}\,\mathrm{d}t$.

Second, our parametric approach allows for a straightforward distinction between *systematic* errors (measured by the difference between the actual and the estimated distributions) and *non-systematic* ones (represented by the usual zero-mean error term which characterizes trial-and-error behavior), whereas the non-parametric approach automatically associates errors (measured by the difference between actual and estimated values) to an "out-of-equilibrium" behavior. For this reason, the non-parametric approach cannot disentangle the two types of errors and is bound to interpret trial-and-error type of behavior as systemic errors in experimental contexts.

3. Inference within the Empirical Framework

3.1. *Characterization of bids and prices.*

In this framework we assume the researcher to only have access to bid data, and that his/her objective is to estimate the unknown parameter vector $\theta = (\theta_1, \theta_2)$ of $F_\theta(\cdot)$. We consider the following two situations which an empirical researcher typically has to deal with.

WINBIDS: The researcher only has information about winning bids, as it would be the case when studying Dutch (descending-price) auction data. The expected winning bid b_t^W in auction *t* conditional on θ and *r* is then given by the following equation:

$$m_W(\theta_1, \theta_2, r) = E(b_t^W | \theta, r) = \int_0^1 \varphi(\theta, s, r)\, dF_\theta(s)^n. \tag{5}$$

ALLBIDS: The researcher has information on *all* submitted bids b_{it}, as it would be the case when studying First-Price auction data. The expected value of bidder *i*'s bid in auction *t* conditional on θ and *r* then takes the following expression:

$$m_A(\theta_1, \theta_2, r) = E(b_{it} | \theta, r) = \int_0^1 \varphi(\theta, s, r) dF_\theta(s). \tag{6}$$

3.2. *Conditions for identification.*

Using data on all bids or on winning bids only, a first question is whether the specification of the bids' first order moments, i.e., equations (5) or (6), is enough to obtain an estimate of the parameters of interest: r and θ. The literature has discussed some conditions for the identification of these parameters in structural first-price auction models. Perrigne and Vuong (2000) define conditions to be fulfilled by the distribution of bids to ensure that a consistent estimate of the distribution of values can be recovered from bids alone when assuming risk neutrality, i.e., $r = 1$. These conditions essentially require $F_\theta(\cdot)$ to behave smoothly at the tails. Campo *et al.* (2002) further show that when assuming risk aversion, some additional restrictions are needed both on the distribution of bids and on the bidders' utility function to ensure identification. In either case, the identification is non-parametric in the sense that it is not restricted to a particular parametric family of distribution of valuations. Within our parametric setting, we also need to impose some conditions to ensure identification of the parameters $(\theta_1, \theta_2) \in \mathbb{R}^2_{++}$ and r.

Proposition 1: *The equilibrium bidding function is not a one-to-one mapping of the parametric space* Θ *for* (θ_1, θ_2, r)*, i.e., there exist infinitely many vectors* (θ_1, θ_2, r) *and* $(\theta_1^*, \theta_2^*, r^*)$ *in* Θ *with* $(\theta_1, \theta_2, r) \neq (\theta_1^*, \theta_2^*, r^*)$ *such that for any* $v \in [0,1]$,

$$\varphi(v; \theta_1, \theta_2, r) = \varphi(v; \theta_1^*, \theta_2^*, r^*).$$

Proof: See the Appendix.

Proposition 1 states that, in general, bids alone do not provide enough information about the parameters of interest since the observed bidding behavior may be explained by at least two different vectors of parameters. Therefore it is immediate to conclude

Claim 1: *The distribution of valuations* $F_\theta(\cdot)$ *and the CRRA parameter r are not identified from information on bids alone*.

Thus, identification requires a restriction on the parametric space Θ. Notice that from (the proof of) Proposition 1, we can also conclude

Claim 2: *For any* $r > 0$, $\varphi(v;1,1,r) = \varphi(v;1/r,1,1)$.

Claim 2 implies that if people is risk averse with some CRRA parameter r and uniform valuations, the observed bidding behaviour is indistinguishable from the bids obtained under risk neutrality ($r = 1$) with valuations drawn from a Beta with $\theta_2 = 1$ and θ_1 the inverse of *r*. Thus, within an empirical framework (i.e., where the researcher has only access to bid data) risk neutrality should be imposed for the identification of (θ_1, θ_2), as in Perrigne and Vuong (2000). In other terms, the estimates of (θ_1, θ_2) can be recovered from bid data. However, imposing risk neutrality is not enough if the researcher uses only the bids' first order moment.

Claim 3: *For any* r, θ_1, θ_2, *there exists* θ_1^*, θ_2^* *such that*

$$m_A(\theta_1, \theta_2, r) = m_A(\theta_1^*, \theta_2^*, r), \text{ and}$$
$$m_W(\theta_1, \theta_2, r) = m_W(\theta_1^*, \theta_2^*, r).$$

Proof: See the Appendix.

Claim 3 simply shows that one more restriction is needed when only using the first order moments of bids. Namely, if we impose risk neutrality ($r = 1$), then we need to restrict either parameter θ_1 or θ_2. Alternatively, if we impose $\theta_2 = 1$ then we need to restrict either r or θ_1 since, by Claim 2, we would then have $\theta_1 = 1/r$. With such a restriction the SBNE bid function is linear in values so that we have

Claim 4: *If* $\theta_2 = 1$ *and r=1, the bidding function is linear in valuations, i.e.,*

$$\varphi(v; \theta_1, 1, 1) = \frac{n-1}{n-1+1/\theta_1} v$$

To summarize, within the empirical framework we will use first order moments of WINBIDS and ALLBIDS. In view of Claim 3, we will leave the parameter θ_1 free and restrict both θ_2 and r to be equal to one, so that we focus analysis to the linear bidding case.[5]

[5] The reader may ask why not to use higher order moments that may provide identification of both θ_1 and θ_2 using for instance, the work of Deltas (2004) who finds the second order moments of winning bids in independent private value auctions. The reason for using only first order moments is that we do not require full-identification of the Beta

3.3. *Tests and Procedures.*

We will use bid data which has been obtained in an experimental framework for first-price auction, to do the following:

- Testing for a homogenous bidding behavior. We first check if the bidders participating in a given session display a homogenous behavior by estimating equation (6) with individual dummies variables for the parameter θ_1. We denote by θ_{1i} the θ_1 parameter of bidder *i* (*i* = *2,...4* and check whether $\varphi(v;\theta_{1i},1,1) = \varphi(v;\theta_{1j},1,1)$ for all $i \neq j$ by defining dummies which take value one if bid corresponds to individual *i* and specify $\theta_{1i} = \theta_1 + \sum_{j=2}^{4} \delta_i 1_{\{i=j\}}$ for $i > 1$. Homogenous bidding implies $\delta_i = 0$ for all $i > 1$. We refer to the null hypothesis of this test as $H_0^{\delta=0}$. Note that rejecting the null of a homogenous behavior in either of our frameworks cannot reasonably mean that bidders are acting in some asymmetric equilibrium since it is common knowledge to them that their respective values are all drawn from the same probability distribution; a condition that leads (almost certainly) to a unique equilibrium bidding strategy.[6] Also, bidders cannot be said to be asymmetric in their risk attitudes, as in Li and Riley (1999), because the equilibrium predictions then assume that the bidders' risk aversion parameters are common knowledge, which is certainly not the case here.[7] For these reasons, rejecting the null of a homogenous behavior in our setting is tantamount to reject the SBNE model as symmetric individuals are found not to be using the same bid function.
- Testing the informational equivalence of bids and prices. The distinction between the ALLBIDS and WINBIDS frameworks allows one to test the appropriateness of the SBNE model to organize the observed behavior. Notice first that if individuals are homogenous and bid in equilibrium, then the θ_1 estimates in ALLBIDS and WINBIDS (i.e., using (6) and (5), respectively) are consistent. Thus, by virtue of consistency, the estimates under ALLBIDS and WINBIDS should converge to the same value as the sample of observations increases. However, for small data sets such as those used in empirical or experimental studies, we

distribution in our framework that involves uniformly drawn valuations and thus, linear equilibrium bidding strategies. The explicit testing of the SBNE will be adressed in the experimental section.

[6] Maskin and Riley (2003) show that unicity is guaranteed if the distribution of values is defined on a finite support and if there is a positive mass at the lower endpoint of support which would have been the case here if only the experiments were conducted with a very small reserve price, e.g., $p^* = 0.01$.

[7] See Campo (2002) for an empirical study of the latter case.

expect differences to be observed simply because the winning bid is only the maximum statistic of the sample of all bids $\{b_1, b_2, ..., b_n\}$ so that it contains less information than $\{b_1, b_2, ..., b_n\}$ about the parameter vector (θ_1, θ_2, r). As the researcher has less information in WINBIDS than in ALLBIDS, the estimates are expected to be more efficient in the latter case: both estimation frameworks should yield similar estimates but with smaller variance in ALLBIDS than in WINBIDS. To this extent, testing the equality of ALLBIDS and WINBIDS point estimates and the stochastically larger variances of WINBIDS of the latter determines if the SBNE model is adequate to explain the observed behavior.[8] Also, since the equilibrium bid function is monotonically increasing in values, the vector of observed bids $\{b_1, b_2, ..., b_n\}$ contains as much information as the vector of value realizations $\left\{v_i = b^{-1}(b_1), i = 1, ...n\right\}$ regarding the estimation of (θ_1, θ_2, r).[9] To this extent, we further conjecture that any difference between the estimates obtained from (6) and those obtained from an environment in which bidders' private values are observed, must asymptotically disappear if bidders play in equilibrium. This last conjecture provides an easy way to assess the effect of knowing the bidders' private values on inferences.

- <u>Implementing Optimal Reserve Prices</u>: The structural analysis of auction data is principally motivated by normative reasons like the design of market institutions and the setting of optimal reserve prices. Once estimates of the bidding function are obtained, equation (4) can be used to analyze how the different informational content of bids and prices affects the determination of optimal reserve prices and thus the seller's expected revenue. Another issue that relates to the implementation of optimal reserve prices is the assumption of a risk averse bidding behavior. Given the identification conditions discussed in the section 3.2, determining optimal reserve prices is impossible to achieve without some prior assumption on the bidders' risk preferences (or an exogenous assessment of bidders' risk attitudes). To evaluate the possible effects of assuming a risk averse behavior on the seller's revenues when optimal reserve prices are to be set, we therefore consider the following two scenarios. We

[8] Clearly, we cannot build a Haussman-type of test for this necessary condition because the estimator would be biased both in ALLBIDS and WINBIDS.

[9] This is because if a statistic S is sufficient for a parameter vector (θ_1, θ_2, r) and $h(\cdot)$ is bijective, then $h(S)$ is also sufficient (see Gourieroux and Monfort, 1995). In our case, as the inverse bid function $v_i = b^{-1}(b_i)$ is bijective it meets this condition. This however is not the case for the relationship between the winning bid b_t^W and the bid sample $\{b_1, b_2, ..., b_n\}$.

assume that bidders are either *i)* risk neutral with values distributed according to a Beta distribution (henceforth, the BETA scenario) or *ii)* constant relative risk averse with uniformly drawn values (henceforth, the UNI scenario).

4. Inference within the Experimental Framework

4.1. *Data Generating Process for bids conditional on valuations.*

When conducting auction experiments, the researcher has full information about the number of bidders, the distribution(s) of values and the bidders' value realizations. In most experiments the distribution of values is chosen to be a uniform on $[\underline{v};\overline{v}]$ so that by normalizing values (and bids) to the $[0,1]$ interval, (3) yields the following SBNE bid function, which is linear in values:[10]

$$b_{it} \equiv \varphi(v_{it},1,1,r) = \frac{n-1}{n-1+r} v_{it} \,. \tag{7}$$

Our approach consists in encompassing this linear model into a general parametric model that allows for non-linear (monotone increasing) bid functions by using the fact that a uniform distribution defined on $[0,1]$ is a particular Beta distribution with $\theta_1 = \theta_2 = 1$. Thus, as shown in claim 4, testing that individuals bid as in (7) is equivalent to conduct a parametric test over θ_1 and θ_2 that checks that the estimated Beta distribution is indeed uniform.[11]

Thus, we will consider the following equation for the bids in an experimental framework:

$$b_{it} = v_{it} - \int_0^{v_{it}} \left[\frac{F_\theta(s)}{F_\theta(v_{it})} \right]^{\frac{n-1}{r}} ds + w_{it} = \varphi(v_{it},\theta,r) + w_{it} \,, \tag{8}$$

[10] We normalize values v_{it} and bids b_{it} as $v_{it} = (\mathrm{v}_{it} - \underline{v})/(\overline{v} - \underline{v})$ and $b_{it} = (\mathrm{b}_{it} - \underline{v})/(\overline{v} - \underline{v})$, respectively.

[11] In the Results section of the paper, we will also report on the case where the estimated bid functions can only be linear in values, i.e., when θ_2 is set equal to 1.

where w_{it} stands for a non-systematic error term with $E(w_{it}\,|\,v_{it}) = 0$. This random term is a reduced-form error that captures deviations from equilibrium behavior that are not explained by $\varphi(v_{it}, \theta, r)$.[12]

We know from Proposition 1 that the true parameters r_0 and $\theta^0 = (\theta_1^0, \theta_2^0)$ cannot be identified from (8) simultaneously. According to this proposition, if bidders are risk averse and have their values drawn from a uniform distribution, then in the SBNE, their behavior would not be distinguishable from the behavior of risk neutral bidders with non-uniform private values. Proposition 1 provides however a condition on the parameters θ_1, θ_2 and *r* for the model to be tested with experimental data. This condition is that if we cannot reject the null that $\theta_2 = 1$, then we have $r = 1/\theta_1$ and the estimated bid function is linear in valuations with a slope equal to $(n-1)/(n-1+\dfrac{1}{\theta_1})$, as predicted by the model for uniformly drawn values. In this framework, non-linear bid functions are characterized by $\theta_2 \neq 1$ and are interpreted as evidence against the SBNE model. Therefore by imposing $r = 1$ in (8), we get:

$$b_{it} = v_{it} - \int_0^{v_{it}} \left[\frac{F_{\theta^0}(s)}{F_{\theta^0}(v_{it})} \right]^{n-1} ds + w_{it}\,. \tag{9}$$

Equation (9) can now be estimated using standard non-linear least squares where the integral must be computed using quadrature numerical procedures.

4.2. *A test of the null of Symmetric Bayes-Nash Equilibrium.*

Restricting *r*=1 not only ensures identification of the θ vector parameter but implies that if the SBNE model holds, then it must be the case that the bidding function is linear (i.e., $\theta_2 = 1$) and θ_1 equals $1/r$. Therefore, we can assess the SBNE model by testing

[12] Notice that while (8) allows for trial-and-error bidding, it does not further assume bidders to stochastically best-respond to their rivals' errors, as QRE models do. To this extent, (8) imposes less structure on the bidders' rationality than QRE type of models.

$$H_0 : \begin{pmatrix} \theta_1^0 \\ \theta_2^0 \end{pmatrix} = \begin{pmatrix} 1/r_0 \\ 1 \end{pmatrix},$$
$$H_1 : \begin{pmatrix} \theta_1^0 \\ \theta_2^0 \end{pmatrix} \neq \begin{pmatrix} 1/r_0 \\ 1 \end{pmatrix}. \tag{10}$$

The null hypothesis in (10) defines a function $h : \mathbb{R} \to \mathbb{R}^2$ such that if the model holds and r_0 is the true constant relative risk aversion parameter, then there is an implicit restriction on the parameter space of θ which is given by $\theta^0 = h(r_0)$. To test this implicit restriction, we consider the distance between a consistent estimate of θ^0 under both the null and the alternative hypotheses, $\hat{\theta}$, and $\hat{\theta}^0$, which is consistent only under the null, and we test whether this distance is statistically significant. A consistent estimator under the null hypothesis is given by $\hat{\theta}^0 = h(\hat{r})$, where $\hat{r}$ is a consistent estimate of the true risk aversion parameter r_0. This consistent estimator can be found by estimating the linear model $b_{it} = \gamma v_{it} + \varepsilon_{it}$ by OLS, in which case $\hat{r} = (n-1)(1-\hat{\gamma})/\hat{\gamma}$.

Write $\varphi_{it}(\theta)$ short for $v_{it} - \int_0^{v_{it}} \left[\frac{F_{\theta^0}(s)}{F_{\theta^0}(v_{it})} \right]^{n-1} ds$. A consistent estimator $\hat{\theta}$ can be found by NLLS, i.e.,

$$\hat{\theta} = \underset{\theta}{ArgMin} : S_{nT}(\theta) \equiv \sum_{i,t} (b_{it} - \varphi_{it}(\theta))^2, \tag{11}$$

which is an asymptotically normal estimator so that $\sqrt{nT}(\hat{\theta} - \theta^0) \to N(0, J_0^{-1} I_0 J_0^{-1})$, where I_0 is the asymptotic variance of the gradient of the objective function, and J_0 stands for the asymptotic Hessian evaluated at the true parameter value.

Following Gourieroux and Monfort (1995), it can be shown that $\sqrt{nT}(\hat{\theta} - \hat{\theta}^0)$ is asymptotically normal with a variance-covariance matrix given by $V = M_h J_0^{-1} I_0 J_0^{-1} M_h^{'}$, where M_h is the orthogonal projection of $\partial h(r_0)/\partial r$ on the space spanned by the columns of J_0. Notice that since the only effective restriction in (10) is the one affecting θ_2^0, the 2×2 matrix V is singular.

Therefore, an asymptotic test of (10) at the required confidence level can be computed using the fact that

$$nT(\hat{\theta} - \hat{\theta}^0)' V^{-}(\hat{\theta} - \hat{\theta}^0) \to \chi^2_{k-q} \,,$$

where V^{-} denotes any generalized inverse of V. If the variance matrix Ω of the non-systematic error term w_{it} is assumed to be a scalar matrix, I_0 can be consistently estimated using the outer product of the gradients $\sum_{i,t} \partial_\theta \varphi_{it}(\theta) \partial_\theta \varphi_{it}(\theta)'$. In general, however, one would expect some degree of both heteroskedasticity and autocorrelation (both of unknown form) in the error term.[13] Therefore we use a Newey-West estimator of I_0, which we denote by $\hat{I}$.[14]

Let $\hat{J} = \partial^2 S_{nT}\left(\hat{\theta}\right) / \partial^2 \theta$ denote the estimate of the Hessian, then the test of the SBNE model is computed as follows:

$$nT(\hat{\theta} - \hat{\theta}_0)'(M_h \hat{J}^{-1} \hat{I} \hat{J}^{-1} M_h')^{+}(\hat{\theta} - \hat{\theta}^0) \to \chi^2_{k-q} \,, \tag{12}$$

with + denoting the Moore-Penrose generalized inverse. The next proposition states that this test reduces to a standard normal test of the significance of the second log-parameter of the estimated Beta distribution, which significantly simplifies the computation of the test in (10):

Proposition 2: Call $\hat{\tilde{\theta}}_2$ *the estimate of the log-parameter* $\tilde{\theta}_2$ *and* $\hat{\sigma}_{\hat{\tilde{\theta}}_2}$ *the estimate of its standard deviation based on the asymptotic distribution. Under the null hypothesis in (10), the t-statistic* $t = \hat{\tilde{\theta}}_2 \big/ \hat{\sigma}_{\hat{\tilde{\theta}}_2}$ *is asymptotically normal with zero mean and unit standard deviation.*

Proof: See the Appendix.

---

[13] Heteroskedasticity may be due to the fact that bids are typically less dispersed at low values than at high values. Autocorrelation instead may be due to the bidders' possible trial-and-error behavior which can induce time-dependence in the residuals w_{it}.

[14] If the only source of heteroskedasticity was due to an increasing dispersion of bids with respect to values, then an alternative approach would be to use Generalized Least Squares (which also assumes non-autocorrelated residuals). We conducted such GLS estimations and found no significant change in our conclusions to report.

4.3. *Testing the Symmetric Bayes-Nash model with Heterogeneous risk attitudes*

So far the SBNE model assumed bidders with homogenous constant relative risk averse preferences. The model proposed by Cox et al. (1982) assumes instead that the bidders' utilities are defined as $u_i(w) = w^{r_i}$, where r_i is i.i.d. according to $G(\cdot)$ which has a positive continuous density $g(\cdot)$ on $(0, r_{\max}]$, and where $r_{\max}$ stands for the risk parameter of the least risk averse bidder. As behavior at first-price auction experiments is typically characterized by a significant overbidding, risk neutrality is often considered as a natural lower bound on bidders' risk aversion so we assume $r_{\max} = 1$. Notice that the additional parameters $G(\cdot), g(\cdot)$ and $(0, r_{\max}]$ are assumed to be common knowledge, and that although bidders have heterogeneous risk parameters, they are still ex-ante symmetric.

In what follows, we use two properties of this model to construct a structural test. A first property is that if bidders' valuations are uniformly distributed, then the individual equilibrium strategies are linear up to the maximum possible bid, b^*, of the least risk averse bidder and they have no closed-form solution for bids greater than b^*. For uniformly drawn valuations on the unit interval, Cox *et al.* (1988) show that $b^* = (n-1)/(n-1+r_{\max})$ and that the equilibrium bid functions have the following expression

$$b_{it} = \varphi(v_{it};1,1,r_i) = \frac{n-1}{n-1+r_i} v_{it} \qquad \text{for } v_{it} \in [0; \varphi^{-1}(b^*(r_{\max});1,1,r_i)].$$

The second property is that if valuations are uniformly distributed then, for a given $r_{\max}$, the distribution $G(\cdot)$ does not affect the linear part of the equilibrium bid functions. This directly follows from the derivation of equilibrium bid functions in Van Boening, Rassenti, Smith (1998). A structural test of this model would thus consist of checking if (1) bidders display a homogenous behavior or not, (2) their bid functions are linear for bids smaller than $b^*(r_{\max})$ and (3) the estimated risk aversion parameters are all smaller or equal to $r_{\max}$.

5. Results

We test our models with the experimental data of Isaac and Walker (1985) which consist of ten sessions of first-price auctions. Each of these sessions involves four different bidders who play for 25 rounds and who have their valuations drawn with replacement from a uniform on $[0,10]$.[15] We conduct our estimations separately for each session, and we assume that any test outcome for a given session s is a Bernoulli variable R_s that takes the value 1 if the null is rejected (at $\alpha = 0.05$, two-tailed) and 0 otherwise, so that $R_s \sim B(1;0.05)$. To draw our conclusions for the ten auction sessions, we run a one-tailed binomial test on the number of rejections given a probability of rejecting the null ("success") of .05, so that $\Re \sim B(10;0.05)$ with $\Re = \sum_{s=1}^{10} R_s$. The null hypothesis of this test is that the probability of wrongly rejecting the null in a given session is $\alpha = 0.05$ and the alternative is that $\alpha > 0.05$, or equivalently that the null tested in each individual session is globally rejected. Thus, the null of this global test is rejected if the probability of observing at least $\Re$ rejections is smaller than $.05$. This simple test allows a straightforward assessment of a given hypothesis by summarizing the evidence in a series of fully independent observations.

5.1. *Empirical Framework*

In this section we assess homogeneity in behavior and the SBNE model with the estimating equations (5) and (6), and we explore the effects of different information conditions on our inferences and on the seller's expected revenues.

- Testing for Symmetric Bayes-Nash Equilibrium behavior

We first check if the bidders participating in a given session display a homogenous behavior by estimating equation (6) with individual dummy variables for the parameter θ_1 and checking whether the δ_{1i} parameters in the specification $\theta_{1i} = \theta_1 + \delta_{1i}$ for $i = 2,3,4$ are jointly non significantly different from zero ($H_0^{\delta=0}$). The statistics in Table 1 indicate that we reject this hypothesis for four sessions in WINBIDS and for one session in ALLBIDS so that according to the binomial test, it is rejected in WINBIDS ($p = .0010$) but not in ALLBIDS ($p = .4013$). Such different outcomes provide some first evidence that behavior is not consistent with the SBNE model.

We proceed with checking the necessary condition for the observed behavior to be a SBNE one, i.e., that for sessions with homogenous bidders, the WINBIDS and ALLBIDS

[15] These data were also analyzed by Cox *et al.* (1988) and Cox and Oaxaca (1996) and Ockenfels and Selten (2001).

frameworks should yield stochastically equivalent θ_1 estimates, with larger standard deviations in WINBIDS than in ALLBIDS. The conduct of a non-parametric randomization test on the statistics of Table 1 indicates that both the estimates and their standard deviations are significantly larger in WINBIDS than in ALLBIDS ($p = .0003$ and $p = .0000$, one-tailed, respectively) so that the symmetric benchmark model for either constant relative risk averse or risk neutral bidders would not be appropriate to recover the unknown distribution of values from this data.

Notice that if we had not conducted these preliminary tests, then we would have concluded, on the basis of Propositions 1 and 2, that the distribution underlying bidders' values is either a Beta distribution with parameters $(\theta_1, 1)$ --- with an average θ_1 estimate of 1.44 in WINBIDS and of 1.19 in ALLBIDS; or a uniform distribution on $[0,1]$, in which case bidders would have a common average risk averse parameter of $1/\theta_1$ (i.e., $r = .84$ in WINBIDS and $r = .69$ in ALLBIDS).

- Implementing Optimal Reserve Prices

The structural analysis of auction data is principally motivated by normative reasons like the design of market institutions and the setting of optimal reserve prices. Table 1 reports the expected revenue increase (in %) for each session following the implementation of an "optimal" reserve price. While the figures indicate that expected revenues increase in both WINBIDS and ALLBIDS when optimal reserve prices are implemented, they remain well below the 2.08% predicted by the SBNE benchmark for risk neutral bidders, and on average, they are greater in ALLBIDS than in WINBIDS (i.e., 1.40% versus .90%, respectively). The lower-than-expected predictions are most likely due to the bidders' tendency to overbid, which skews the estimated distributions towards higher valuations on the unit interval. Indeed, since the θ_1 estimates are all greater than 1, the mean and standard deviation of the estimated Beta distribution are respectively higher and lower than those of a uniform defined on [0,1]. With such right-skewed distributions, the implementation of optimal reserve prices extracts less of buyers' information rents than if values were uniformly distributed, so that the seller's expected revenues increase by less. Given our previous finding that the WINBIDS estimates are significantly greater than the ALLBIDS ones, this also explains the lower predictions in WINBIDS than in ALLBIDS.

Another issue that relates to the implementation of optimal reserve prices is the role of risk aversion in bidders' preferences (Riley and Samuelson, 1981, Proposition 5). In view of the non-identification results of Propositions 1 and 2, determining optimal reserve prices from bid data is impossible to achieve without some prior assumption on bidders' preferences or an exogenous assessment of bidders' attitudes towards risk. For this reason, and to evaluate the possible effects of risk aversion on the seller's revenues when optimal reserve prices are to be set, we consider the following two scenarios. We assume that bidders are either *i)* risk neutral with values distributed according to a Beta distribution (henceforth, the BETA scenario) or *ii)* constant relative risk averse with uniformly drawn values (henceforth, the UNI scenario).

Figure 1 shows how expected revenues relate to reserve prices for three different degrees of risk aversion, the first two of which corresponding to the average estimates for the WINBIDS and ALLBIDS cases (cf. Panels A and B; the markers represent optimal reserve prices). In panel A, expected revenues (without reserve prices) are equal to $.6456$ if the BETA scenario applies with $\theta = (1.19, 1)$ and to $.6249$ if the UNI scenario applies with $r = 1/1.19 = .84$. Once "optimal" reserve prices are determined and implemented, expected revenues rise to $.6545$ and $.6289$, respectively, leading to increases of 1.38% and .63%. Hence, if the exact distribution of the bidders' values is unknown, the seller could be advised on the grounds of expected revenue maximization to assume the BETA scenario for setting optimal reserve prices. The plots of panels B and C, which assume higher degrees of risk aversion ($r = .69$ and $.50$, respectively), further indicate that the difference in expected revenues between the two scenarios increases with risk aversion and that the benefits from implementing optimal reserve prices in UNI vanish with risk aversion. For example, when $r = .50$, a mid-range figure of estimates reported by several empirical and experimental studies, the optimal reserve price would be equal to $.023$ and would yield virtually no revenue increase (i.e., the difference being of the order 10^{-6}).

Eventually, the plots of panel C reveal that setting a reserve price that would be optimal in scenario BETA, would also disappoint the seller's revenue expectations by -14.47% if the actual scenario is UNI with $r = .50$ (cf., a in panel C). Implementing such a reserve price would also lead to an expected revenue decrease of -4.63% when compared to setting a reserve price corresponding to risk averse bidders with values that are uniformly drawn (i.e., b in Panel C). This numerical example illustrates how important it is for the seller to know which of the two scenarios applies when a reserve price is to be implemented. One way of deciding which scenario applies best would consist in comparing the observed average revenues to those predicted in each scenario and to cast the one that fits the data best.[16] Figure 2 reports the observed and predicted revenues of each session when no reserve price is implemented. The plots indicate that the BETA scenario fits the data best and that within each scenario, the WINBIDS predictions (empty markers) fit the observed values almost always better than the ALLBIDS predictions (solid markers). The latter, however, is not surprising given that the seller's revenues are determined by winning bids. With finite samples of data, the ALLBIDS predictions can therefore only be more noisy than the WINBIDS ones or, as it is the case here, biased by subjects' out-of-equilibrium play.

To summarize, our comparison suggests that the BETA scenario should be used to determine optimal reserve prices. As we could not reject homogeneity in the ALLBIDS case (at $p = .4013$), inferences and recommendations should also account of all bids to fulfill the model's assumptions, although the WINBIDS case outperforms the ALLBIDS case in terms of goodness-of-fit (cf. Figure 2). Unfortunately, the diagnosis of an out-of-equilibrium play prevents a general assessment of the objectives *choosing the model which basic assumptions are fulfilled by data* versus *choosing the one that fits data best* for the determination of optimal reserve prices. Given this out-of-equilibrium behavior, the recommended reserve price may indeed be too high and create a risk of reducing competition more than necessary, which would decrease the seller's expected revenues. Further experimental investigations would be needed to determine which of these objectives tends to outperform the other in terms of expected revenue predictions.

In what follows, we check to what extent the SBNE model for homogenous risk attitudes, or its augmented version for heterogeneous risk attitudes, explains bidding behavior at first-price auction experiments.

5.2 *Experimental Framework*

- Linear Bid functions: Constrained Estimation

To ensure a transition from an empirical to an experimental framework, we first consider a hybrid setup where the researcher still has limited information about the actual distribution of bidders' values (as in the previous section), but has full information about the bidders' value realizations (as is the case in experiments). The purpose of this setup is to assess the effect of

[16] Such a comparison is possible only if the seller's expected revenue is always greater in one of the scenarios,

knowing the bidders' value realizations on our conclusions when the actual distribution $F_\theta(\cdot)$ is still imperfectly known. Such a framework implies the estimation of a variant of (6) which includes the bidders' value realizations and which can be obtained from (9) by imposing the constraint that $\theta_{2i} = 1, \forall i = 1,...,4$. That is, we restrict *i*'s bid function to be linear with a slope equal to $(n-1)/(n-1+1/\theta_1)$. It is important to note here that because of this restriction, we cannot test the SBNE model for homogenous or heterogeneous risk averse bidders but only assess the goodness-of-fit of a proxy of it.

In what follows, we first check for a homogenous behavior by testing $H_0^{\delta=0}$. For sessions that do not reject homogeneity, we estimate θ_1 in (9) without individual dummies whereas for sessions that reject this hypothesis, we estimate the four individual θ_{1i} parameters of a session after discarding all bid observations greater than or equal to $b^*(r_{\max})$. The estimation and test results are reported in Table 2. They indicate a rejection of homogeneity in eight sessions so that this basic hypothesis of the SBNE model is rejected according to the binomial test ($p = .0000$). The estimated risk parameters for sessions with homogenous or heterogeneous bidders average at $.42$ and are significantly different from those of Cox and Oaxaca (1996) who use the same data and report an average of $.33$ ($p = .0234$ according to a two-tailed Randomization test with 10 paired observations). We attribute this difference to the fact that Cox and Oaxaca (1996) estimate affine bid functions which happen to have significantly negative intercept terms whereas we estimate bid functions without intercept terms to comply with the SBNE model. The presence of negative intercepts, which tend to increase the estimated slopes of linear bid functions, could therefore explain their higher risk estimates. For the two sessions that did not reject homogeneity (Sessions 2 and 3), the average risk parameter is $.44$, which is equal to the one reported by Campo, Perrigne and Vuong (2002) for timber auctions that were conducted with an average of 3.7 bidders, and which is similar to those reported by Kagel, Harstad and Levin (1986) for six-bidder auctions ($r = .49$) and by Goeree *et al.* (2000) and Armantier and Treich (2005) for two-bidder auctions ($r = .48$ and $.39$, respectively).

Hence, although the above studies involved different econometric procedures and/or assumed different equilibrium models of behavior, the relative invariance of risk estimates to the number of participating bidders seems to support the SBNE model and to suggest that the risk aversion hypothesis fits the data well. Notice, however, that if we had not first checked for a

---

which is the case here.

homogenous behavior, we would have found an average $\hat{\theta}_1$ estimate of 2.837, or equivalently $r = .35$. Therefore, conditioning the estimation on the type of behavior observed (i.e., homogenous or not) causes a 36.5% change in the $\hat{\theta}_1$ estimates or a 20% change in the average risk parameter. Furthermore, the difference between these estimates and the empirical ones reported in the previous section suggests that once behavior has been diagnosed "out of equilibrium", structural inferences become highly sensitive to the quality of the information available and in particular, that they do not necessarily improve with the quality of information. We illustrate this point by reporting in Figure 2 the expected revenues associated to the $\hat{\theta}_1$ estimates of Table 2 when assuming a BETA and a UNI scenario. When compared to the ALLBIDS predictions of Figure 2, the UNI scenario now appears to outperform BETA in explaining the data. To this extent, improving the quality of information available, i.e., taking account of the bidders' value realizations, does not warrant better estimates or structural inferences.

We proceed with a comparison of our results to those of Bajari and Hortaçsu (2005) who report an aggregate estimate for three- and six-bidders auctions of $r = .27$ when the data are analyzed with Quantal Response Equilibrium (QRE) models, and of $r = .158$ when non-parametrically estimating the SBNE model. These estimates are interesting in two aspects. First, when we estimate an aggregate risk parameter for three- and six-bidders auctions with our parametric approach, we find an estimate of $r = .493$. We attribute such different estimates to the fact that in experimental contexts, non-parametric estimation procedures assimilate trial-and-error behavior to "systematic errors" and are thus very likely to interfere with ones conclusions, especially if the observed behavior is out-of-equilibrium. A comparison of the estimates reported in Table 1, which were obtained in an empirical framework that does not account for trial-and-error behavior, to those reported in Table 2, which have been corrected for such behavior illustrates the possible variability in the estimates that one can encounter within a parametric estimation framework. Second, when we separately estimate risk aversion parameters for three- and six-bidders auctions, we find an aggragte estimate of $r = .43$ when $n = 3$ and of $r = .72$ when $n = 6$. Since the Dyer *et al.*'s experiments were designed such the same bidders submitted bids for markets with $n = 3$ and for markets with $n = 6$, these estimates refer to the same set of bidders and are clearly *not* invariant to the number of bidders, as claimed by theory; therefore they strongly suggest an out-of-equilibrium behavior.

- Non-Linear Bid functions: Unconstrained Estimation

We now estimate (9) with individual dummies for *both* parameters so that we allow for non-linear (monotone increasing) bid functions. We will use essentially the same testing procedure as the one just described: for each session, we first check for a homogenous behavior by testing $H_0^{\delta=0}$. If we cannot reject this hypothesis, we estimate θ_1 and θ_2 in (9) without individual dummies and test the SBNE model by checking that the common θ_2 parameter estimate is equal to 1. If we do reject homogeneity, we check for linearity in the estimated bid functions by testing that the θ_{2i} estimates are jointly equal to 1. Finally, we test the SBNE model for heterogeneous risk attitudes by testing that the θ_{2i} estimates are still jointly equal to 1 after having discarded all bid observations greater than or equal to $b^*(r_{\max})$, and by checking that the estimated risk parameters are all smaller or equal to $r_{\max}$. If neither of these hypotheses is rejected for a particular session, then we do not reject the SBNE model for that session. Our conclusion for the ten sessions is again based on the results of a binomial test, as previously outlined.

The estimation and test results are reported in Table 3. A comparison of the predictive accuracy of the models estimated in Tables 2 and 3 indicates that allowing for non-linear bid functions considerably improves the model's adjusted goodness-of-fit. The test results in Table 3 also indicate a rejection of homogeneity in nine sessions ($p = .0000$, according to the binomial test), so that we focus analysis on the SBNE model for heterogeneous risk averse bidders. The $H_0^{\theta_{2i}^*=1}$ test statistics indicate that we reject this model in seven sessions so that we reject it as a possible explanation of the observed behavior ($p = .0000$, according to the binomial test).[17]

The average coefficient for the three sessions that did not reject the model is equal to $.37$ which is more in line with the average of $.33$ reported by Cox and Oaxaca (1996), but lower than the one reported by Chen and Plott (1998) ($r = .52$) for auctions with three bidders who had their values drawn from identical but non-uniform distributions. This difference, however, could be due to the different assumptions regarding the distribution $G(\cdot)$ of risk parameters and the value of $r_{\max}$. Chen and Plott (1998) assume individual risk parameters to be LogNormally

[17] To test the model we had to discard 110 observations with $b > b^*(r_{\max} = 1)$ from analysis, which is less than half the expected number of observations that we would have to discard, had we truncated the data at $v_i^* = .75$, as in Kagel and Levin (1985). See Cox, Smith and Walker (1992, pp 1399-1400) for a discussion on this truncation issue. As all bidders who display linear bid functions on $[0, b^*]$ also have $r_i \leq 1$ at $\alpha = .05$, we do not report the outcomes of this test.

distributed on $[0;+\infty)$, thereby allowing for infinitely risk loving bidders, whereas we assumed $r_{\max} = 1$. When we allow for risk loving bidders and conduct our estimations with $r_{\max} = 2$, the statistics in Table 3 indicate that we still reject the model at $p = .0010$ according to the binomial test.[18]

Figure 3 reports individual bid data and the estimated bid functions for each of the ten sessions. These plots usually indicate a less-than-proportional bidding at high valuations, which can be consistent with the SBNE model if it occurs at bids greater than $b^*(r_{\max})$. We find a similar pattern in the data used by Bajari and Hortaçsu (which rejected the SBNE model for homogenous or heterogeneous risk attitudes, according to the binomial test), and in those of Cox *et al.* (1982) for auctions with 3 and 5 bidders (which also rejected the models tested). It is also present in auctions with 6 or 10 bidders as Kagel and Roth (1992) report relative deviations from risk neutral SBNE bids that are negatively and significantly correlated to the bidders' values. Battigali and Siniscalchi (2003) show that such a behavior is rationalizable in terms of bidders' beliefs even if they have risk neutral preferences. Their argument, however, produces a sizeable less-than-proportional bidding in auctions with two bidders and uniformly drawn values, but it becomes hardly noticeable when there are four bidders or more.

Another possible equilibrium explanation is that bidders misperceive their winning probabilities: the less-than proportional bidding at high values would then witness their overconfidence in winning upon receiving high values; a well-documented phenomenon in decision-making theory and psychology (Prelec, 1998). Goeree *et al.* (2002) allow for such a probability misperception in a QRE model and find that it explains the data almost as well as risk aversion. Armantier and Treich (2005) adapted the experimental setup of Goeree, *et al.* to elicit subjects' winning probabilities and disentangle the effect of risk aversion from the one of probability misperception. They report a significant underestimation of winning probabilities leading to overbidding over the whole range of values and show that accounting for both probability misperception and risk aversion improves the models' fit and makes risk aversion less salient ($r = .73$ instead of $.39$).[19] Such findings indicate that within a QRE framework,

[18] When $r_{\max} = 2$, we had to discard 335 observations with bids greater than *.60*, which represent 225 more observations than with $r_{\max} = 1$. Note that the number of observations discarded is expected to increase with $r_{\max}$ and therefore, to reduce the power of any inference based on this approach.

[19] Since our approach is based on an equilibrium condition where the beliefs are not given by the 'true' distribution, it could easily be re-interpreted in terms of misperception hypothesis within a SBNE model with risk neutral bidders. Note however that any such interpretation requires a prior assessment of the validity of the SBNE equilibrium, which we found to be refuted.

probability misperception seriously challenges risk aversion in explaining the overbidding phenomenon. Whether these findings apply to auctions with more than two bidders or within the SBNE model remains to be determined.

Finally, the overall comparison of our results for the empirical and experimental frameworks reveals that information about the bidders' value realizations does not matter as much as knowing the bidders' distribution of values when it comes to assess the SBNE model of bidding with structural methods.

6. Conclusion

We analyzed laboratory data on first-price private value auctions with structural econometric methods pertaining to either empirical or experimental settings. For both settings, the results of our study shed a new light on the symmetric Bayes-Nash equilibrium model to explain first-price bid data. From an empirical standpoint --- when not taking account of the bidders' value realizations and assuming limited knowledge about the bidders' distribution of values --- we use two basic implications of the model as tests for equilibrium behavior. The first comes from the homogeneity of bidding behavior and implies identical bidding strategies, which is fundamental to the determination of the symmetric equilibrium strategy. The second relies on the finite sample properties of estimators and determines if the observed behavior can possibly be in a Symmetric Bayes-Nash Equilibrium, an assumption that is typically imposed in the structural estimation of symmetric auction models.[20] Passing these tests provides some warrant that both predictions and recommendations based on structural econometric models will not be ill-fated. Their non-fulfillment, however, suggests the presence of an out-of-equilibrium behavior. Our study has shown how structural inferences may vary when this is the case and lead to different recommendations, depending on the prior check of the game's fundamental assumptions and on the quality of the information available. Most importantly for field studies, it has shown that structural inferences do not necessarily improve with the quality of information, and may even worsen.

[20] Similar implications can be derived for the structural assessment of asymmetric auction models. The Bayes-Nash equilibrium predictions for asymmetric first-price auctions predict not only how bidders should act in equilibrium but also how each type of bidder should best-respond to the other type(s) (cf. Lemma 3.1 and Proposition 3.3 of Maskin and Riley, 2000, or Propositions 9 and 10 of Li and Riley, 1999). Such properties of the equilibrium could be seen as necessary conditions to be fulfilled by data for the observed behavior to be labeled "in equilibrium".

From an experimental standpoint --- when taking account of all the information available --- our study indicates a rejection of the Bayes-Nash model for either homogenous or heterogeneous constant relative risk averse bidders. The outcomes indicate in particular that although the model can be found to fit the data well (with important differences between parametric and non-parametric estimates), it does not support them (at $\alpha = .005$) when it is tested jointly for all the participating bidders in a session. The observed behavior is usually found to display heterogeneity and an overbidding that is less-than-proportional at high values, both of which remain so far unexplained.

Finally, our examination of the data from both empirical and experimental standpoints with similar estimation procedures also put in perspective the overall effect of knowing the bidders' true distribution of values or of the bidders' value realizations on inferences. It has shown in particular that the main limitation of non-experimental data, i.e., not knowing the bidders' distribution of values, makes it less likely to reject models than when one uses experimental data, especially if the basic underlying assumptions are not checked.

References

Armantier O. (2002), "Deciding between the common and private values paradigm: an application to experimental data", *International Economic Review*, 43(3), 783-801.

Armantier O. and N. Treich (2005), "Overbidding in independent private-values auctions and misperception of probabilities", *mimeo*, University of Montréal.

Athey S. and P.A. Haile (2002), "Identification of standard auction models", *Econometrica*, 70 (6), 2107-2140.

Bajari P. and A. Hortaçsu (2005), "Are structural estimates of auction models reasonable? evidence from experimental data", *Journal of Political Economy*, 113(4),703-741.

Battigali P. and M. Siniscalchi (2003), "Rationalizable bidding in first-price auctions", *Games and Economic Behavior*, 45, 38-72.

Campo (2002), "Attitudes towards risk and asymmetric bidding: evidence from construction procurements", *mimeo*, University of North-Carolina, Working paper.

Campo S., E. Guerre, I. Perrigne and Q. Vuong (2000), "Semiparametric estimation of first-price auctions with risk averse bidders", *mimeo*, University of Southern California.

Campo S., I. Perrigne and Q. Vuong (forthcoming), "Asymmetry in first-price auctions with affiliated private values", *Journal of Applied Econometrics.*

Chen K.-Y. and C.R. Plott (1998), “Nonlinear behavior in sealed bid first price auctions”, *Games and Economic Behavior*, 25, 34-78.

Cox J.C. and R.L. Oaxaca (1996), “Is bidding behavior consistent with bidding theory for private value auctions”, in R. Mark Isaac (ed*.), Research in Experimental Economics*, vol. 6, Greenwich, CT: JAI Press, pp 131-148.

Cox J.C., V.L. Smith and J. Walker (1982), “Auction market theory of heterogeneous bidders”, *Economics Letters*, 9(4), 319-325.

Cox J.C., V.L. Smith and J. Walker (1983), “Tests of a heterogeneous bidder’s theory of first-price auctions”, *Economics Letters*, 12(3-4), 207-212.

Cox J.C., V.L. Smith and J. Walker (1985), “Experimental development of sealed-bid auction theory; calibrating controls for risk aversion”, *American Economic Review (Papers and Proceedings)*, 75(2), 160-165.

Cox J.C., V.L. Smith and J. Walker (1988), “Theory and individual behavior of first-price auctions”, *Journal of Risk and Uncertainty*, 1, 61-99.

Cox J.C., V.L. Smith and J. Walker (1992), “Theory and misbehavior of first-price auctions: comment”, *American Economic Review*, 82(5), 1392-1412.

Deltas G. (2004), “Asymptotic and small sample analysis of the stochastic properties and certainty equivalents of winning bids in independent private values auctions”, *Economic Theory*, 23, 715- 738.

Dyer D., J. Kagel and D. Levin (1989), “Resolving uncertainty about the number of bidders in independent private value auctions: an experimental analysis”, *RAND Journal of Economics*, 20(2), 268-279.

Friedman D. (1992), “Theory and misbehavior of first-price auctions: comment”, *American Economic Review*, 82(5), 1374-1378.

Goeree J.K., C.A. Holt and T.R. Palfrey (2002), “Quantal response equilibrium and overbidding in private-value auctions”, *Journal of Economic Theory*, 104(1), 247-272.

Gouriéroux C. and A. Monfort (1995), Statistics and Econometrics Models, Vol.2, Cambridge University Press.

Guerre E., I. Perrigne and Q. Vuong (2000), “Optimal nonparametric estimation of first price auctions”, *Econometrica*, 68(3): 525-574.

Harrison G.W. (1989), “Theory and misbehavior of first-price auctions”, *American Economic Review*, 79, 749-762.

Harrison G.W. (1990), "Risk attitudes in first-price auction experiments: a bayesian analysis", *Review of Economics and Statistics*, 72(3), 541-546.

Harrison G.W. (1992), "Theory and misbehavior of first-price auctions: reply", *American Economic Review*, 82(5), 1426-1443.

Holt C.A. (1980), "Competitive bidding for contracts under alternative auction procedures", *Journal of Political Economy*, 88(3), 433-445.

Isaac R.M. and J.M. Walker (1985), "Information and conspiracy in sealed bid auctions", *Journal of Economic Behavior and Organization*, 6, 139-159.

Kagel J.H. (1995), "Auctions: a survey of experimental research" in *Handbook of Experimental Economics*, Kagel J.H. and Roth A.E. (eds), Princeton University Press.

Kagel J.H., R.M. Harstad and D. Levin (1987), "Information impact and allocation rules in auctions with affiliated private values: a laboratory study", *Econometrica*, 55(6), 1275-1304.

Kagel J.H. and A.E. Roth (1992), "Theory and misbehavior of first-price auctions: comment", *American Economic Review*, 82(5), 1379-1391.

Laffont J.-J., H. Ossard and Q. Vuong (1995), "Econometrics of first-price auctions", *Econometrica*, 63(4), 953-980.

Laffont J.-J. and Q. Vuong (1993), "Structural econometric analysis of descending auctions", *European Economic Review*, 37, 329-341.

Li H. and J. Riley (1999), "Auction Choice", *mimeo*, UCLA.

Maskin E. and J. Riley (2000), "Asymmetric auctions", *Review of Economic Studies*, 67, 413-438.

Maskin E. and J. Riley (2000), "Equilibrium in sealed high bid auctions", *Review of Economic Studies*, 67, 439-454.

Maskin E. and J. Riley (200?), "Uniqueness and equilibrium in sealed high bid auctions", *Games and Economic Behavior*, ??, ???-???.

Merlo A and A. Schotter (1992), "Theory and misbehavior of first-price auctions: comment", *American Economic Review*, 82(5), 1413-1425.

Ockenfels A. and R. Selten (2005), "Impulse balance equilibrium and feedback in first-price auction", *Games and Economic Behavior*, 51(1), 155-170.

Neugebaeur T. and R. Selten (2006), "Individual behaviour of first-price sealed auctions: the importance of information feedback in experimental markets", *Games and Economic Behavior*, 54(1), 183-204.

Paarsch H.J. and H. Hong (2006), An Introduction to the Structural Econometrics of Auction Data, *MIT Press*.

Paarsch H.J. (1992), "Deciding between the common and private value paradigms in empirical models of auctions", *Journal of Econometrics*, 51, 191-215.

Perrigne I. and Q. Vuong, (2000), "Structural Econometrics of First-Price Auctions: A Survey of Methods", *Canadian Journal of Agricultural Economics*, 47, 203-223.

Prelec D. (2006), "The probability weighting function", *Econometrica*, 66, 497-527.

Riley J.G. and W.F. Samuelson (1981), "Optimal auctions", *American Economic Review*, 71(3), 381-392.

Van Boening M.V., S.J. Rassenti and V.L. Smith (1998), "Numerical computation of equilibrium bid functions in a first-price auction with heterogeneous risk attitudes", *Experimental Economics*, 1, 147-159.

Vickrey W. (1961), "Counterspeculation, auctions, and competitive sealed tenders", *Journal of Finance*, 16, 8-37.

Walker J.M., V.L. Smith and J.C. Cox (1990), "Inducing risk-neutral preferences: an examination in a controlled market environment", *Journal of Risk and Uncertainty*, 3, 5-24.

TABLES AND FIGURES

Table 1: Empirical framework

Session	WINBIDS $H_0^{\delta=0}$ [a]	$\hat{\theta}_1$	s.d.[b]	p^{*c}	ΔER[d]	ALLBIDS $H_0^{\delta=0}$	$\hat{\theta}_1$	s.d.	p^*	ΔER
1	.8115	1.42	.1253	.5367	0.89	.7199	1.14	.0876	.5131	1.53
2	.0179	1.54	.1175	.5459	0.72	.9773	1.24	.0831	.5218	1.25
3	.1251	1.17	.1170	.5157	1.43	.6472	1.11	.0787	.5103	1.63
4	.0121	1.48	.1223	.5414	0.80	.5732	1.25	.0836	.5227	1.22
5	.5596	1.50	.1160	.5429	0.77	.7536	1.25	.0849	.5227	1.22
6	.1309	1.38	.1197	.5335	0.95	.4876	1.11	.0846	.5103	1.63
7	.0855	1.80	.1619	.5644	0.47	.6663	1.29	.0892	.5261	1.13
8	.0040	1.19	.1321	.5175	1.38	.0339	1.05	.0865	.5048	1.86
9	.0007	1.51	.1438	.5436	0.75	.9235	1.20	.0826	.5184	1.35
10	.0946	1.44	.1335	.5382	0.85	.5310	1.26	.0840	.5236	1.20

[a]: Rejection probability; H_0: $\theta_{1i} = \theta_{1j}, \forall i \neq j$; [b]: Newey-West standard deviations; [c]: Optimal reserve price assuming $B(\hat{\theta}_1, 1)$ and $r = 1$; [d]: % change in Expected Revenues when p^* is implemented $\Delta ER = 100 \times [ER(p^*) - ER(p^* = 0)] / ER(p^* = 0)$.

Table 2: Experimental framework: Constrained estimation

Session	$\hat{\theta}_1$	s.d.	p^*	$H_0^{\delta=0}$	$\hat{\theta}_1^{*}$ [a]	r [b]	AIC[c]
1	2.47	.0779	.6043	.0000	2.43	[.42]	-109.87
2	3.08	.2942	.6335	.2347	5.74	.32	-250.38
3	1.78	.1366	.5630	.1944	2.14	.57	-167.76
4	2.93	.2701	.6268	.0384	5.03	[.37]	-229.02
5	3.49	.1318	.6503	.0000	6.42	[.29]	-27.46
6	1.92	.1273	.5723	.0000	2.04	[.54]	-202.94
7	5.50	.2208	.7115	.0000	5.35	[.20]	-129.31
8	1.51	.1825	.5436	.0000	1.64	[.75]	-190.03
9	2.82	.1276	.6217	.0000	2.88	[.36]	-31.32
10	2.87	.1828	.6241	.0000	5.45	[.35]	-216.43

[a]: Parameter estimated after truncation of bids greater than b; [b]: Average estimated risk averse parameter $r = 1/\hat{\theta}_1$ if bidders are homogenous and $[r] = \sum_{i=1}^{4} 1/4\hat{\theta}_{1i}^{*}$ if they are heterogeneous; [c]: Akaike Information Criteria.

Table 3: Experimental framework: Unconstrained estimation

Sess.	$\hat{\theta}_1$	$\hat{\theta}_2$	$H_0^{\delta=0}$ [a]	$H_0^{\theta_2=1}$ [b]	$H_0^{\theta_{2i}=1}$ [c]	$H_0^{\theta^*_{2i}=1}$ [d] (r_{max} = 1)	r [e] (r_{max} = 1)	AIC	$H_0^{\theta^*_{2i}=1}$ (r_{max} = 2)	r (r_{max} = 2)
1	1.98	0.61	.0521	.1076	.4340	.9999	.62	-99.58	.9999	.66
2	15.19	6.89	.0000	.0185	.0000ø	.0000 ø	---	168.04	.9998 ø	[.13]
3	4.01	2.91	.0127	.0000	.0000	.0000	---	-136.23	.0000	---
4	10.23	5.09	.0096	.0007	.0000	.0000	---	-207.21	.0000	---
5	12.79	5.49	.0000	.0000	.0000	.0000	---	67.49	.9637	[.18]
6	2.57	1.56	.0127	.1571	.0013	.0064	---	-188.53	.9999	[.41]
7	4.81	0.71	.0000	.6095	.9198 ø	.9828 ø	[.22]	95.89	.2560 ø	[.17]
8	1.81	1.32	.0000	.3251	.0002 ø	.0008 ø	---	-95.75	.9177ø	[.80]
9	3.02	1.13	.0000	.7513	.0005	.3727	[.28]	-8.02	.0000	---
10	10.43	5.22	.0130	.0000	.0099	.0947	---	-197.02	.1587	[.16]

Note: We discarded outliers from the data of Sessions 2, 7 and 8 (ø) to allow the estimation procedure to converge.

[a]: Rejection probabilities; $H_0^{\delta=0}: \theta_{1i} = \theta_{1j}$ and $\theta_{2i} = \theta_{2j}, \forall i \neq j$; [b]: $H_0^{\theta_2=1}: \theta_2 = 1$; [c]: $H_0^{\theta_{2i}=1}: \theta_{2i} = 1, \forall i$;

[d]: $H_0^{\theta^*_{2i}=1}: \theta_{2i^*} = 1, \forall i$ (assuming $r_{max} = 1$); [e]: Average estimated risk parameter for sessions that did not reject $H_0^{\delta=0}$ [or $H_0^{\theta^*_{2i}=1}$].

Figure 1: Expected revenues, optimal reserve prices and risk aversion

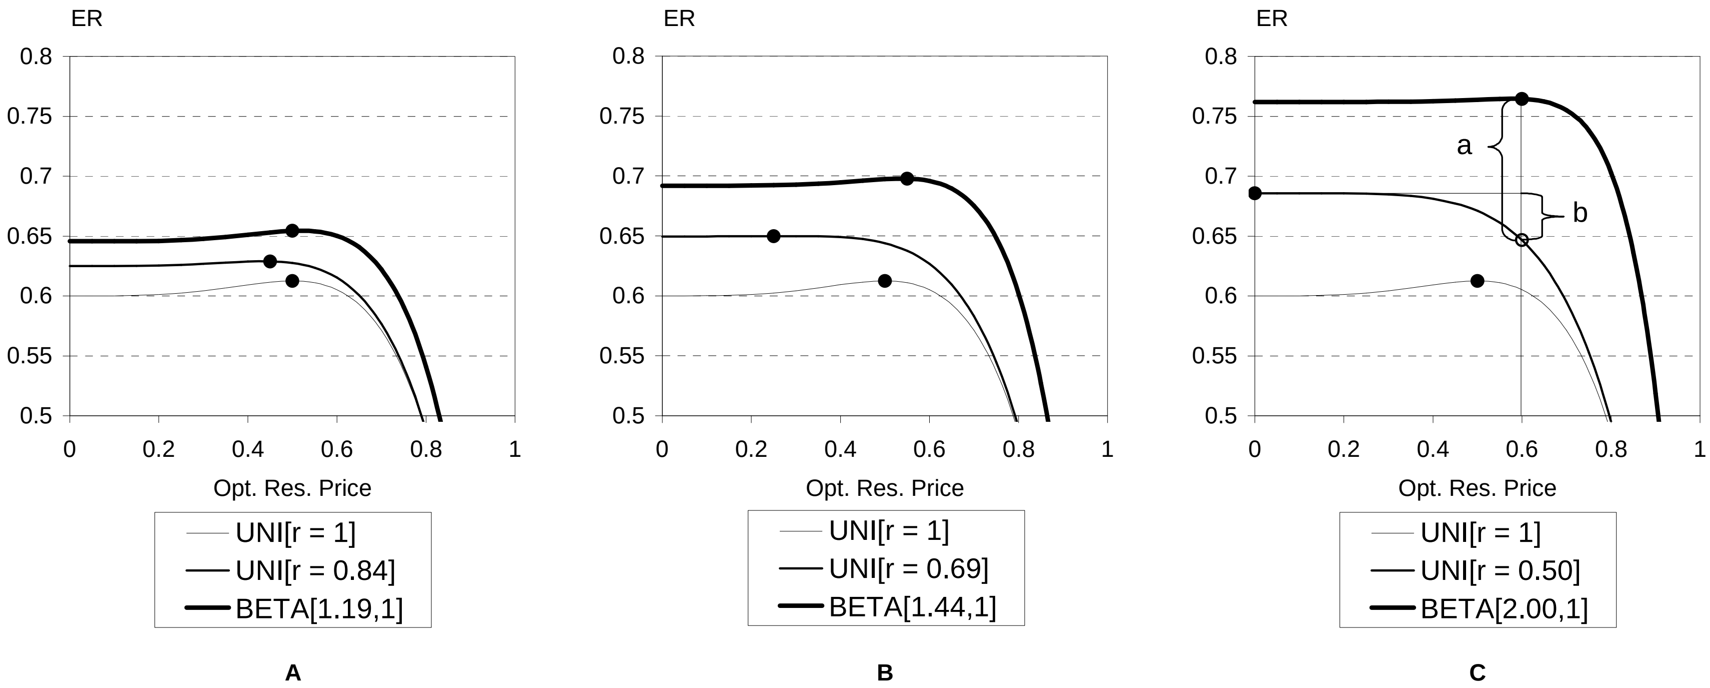

Note: Panels A, B and C report the seller's expected revenues corresponding to different reserve prices, and different assumptions about the bidders' distributions of values (UNI: Uniform, BETA: Beta) and risk preferences (r). Filled dot markers stand for optimal reserve prices (given the distribution of values and risk preferences). In Panel C, a stands for the decrease in expected revenues when setting a reserve price that is optimal for risk neutral bidders with values drawn from a Beta distribution B(2,1) whereas bidders are risk averse (with $r = 0.5$) and have their values drawn from a uniform distribution on [0,1]. b measures the difference in expected revenues between the case where the wrong optimal reserve price has been implemented (i.e., for risk neutral bidders with beta distributed values) and the case where a reserve price for risk averse bidders with values drawn from a uniform has been implemented.

Figure 2: Observed and expected revenues

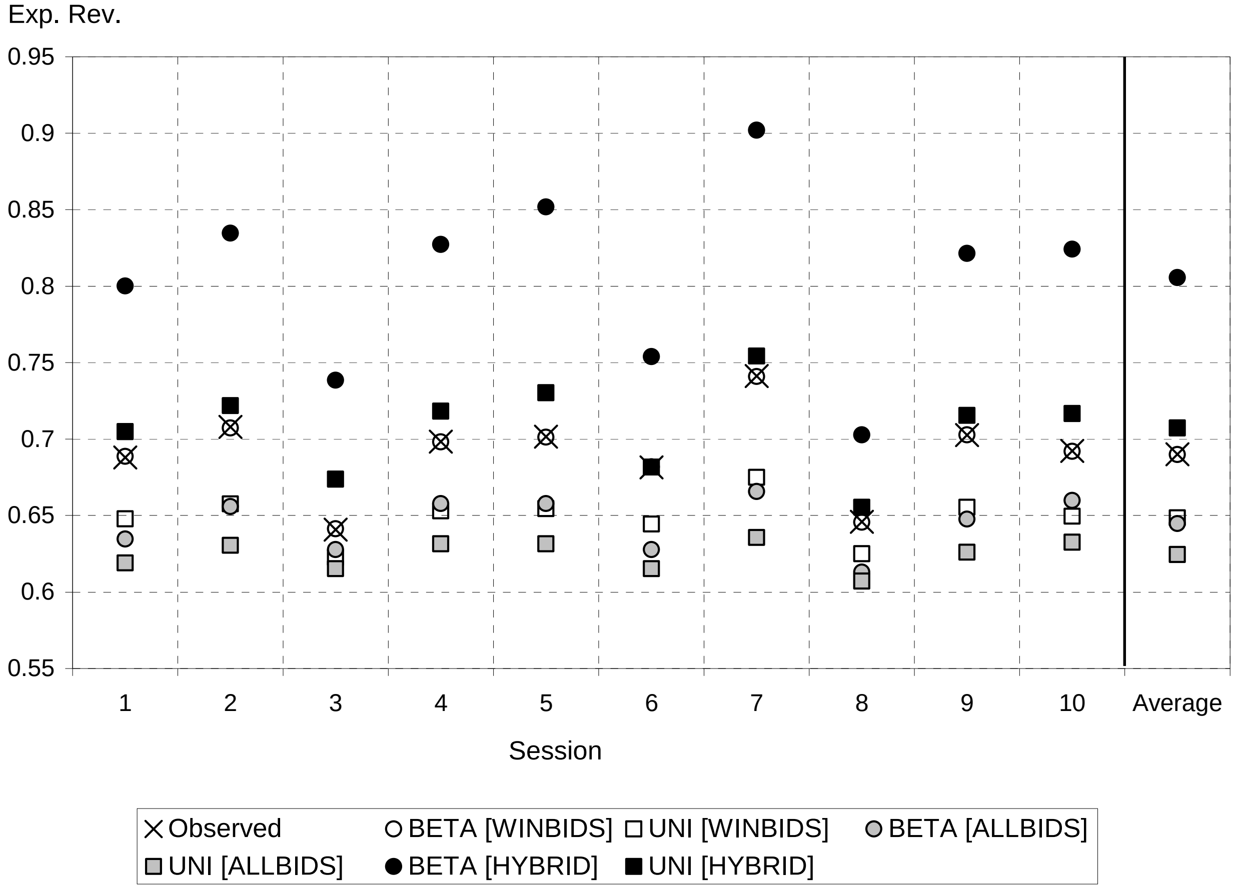

Note: The figure reports observed and expected revenues under various conditions. BETA[·] assumes risk neutrality and values drawn from a Beta distribution (with the parameter estimates of Table 1 or 2). UNI [·] assumes risk aversion and values drawn from a uniform distribution (with the parameter estimates of Table 1 or 2). WINBIDS: estimates use information on winning bids. ALLBIDS: estimates use information on all bids. HYBRID: estimates use information on all bids and on value realizations.

Figure 3: Estimated Individual Bid Functions

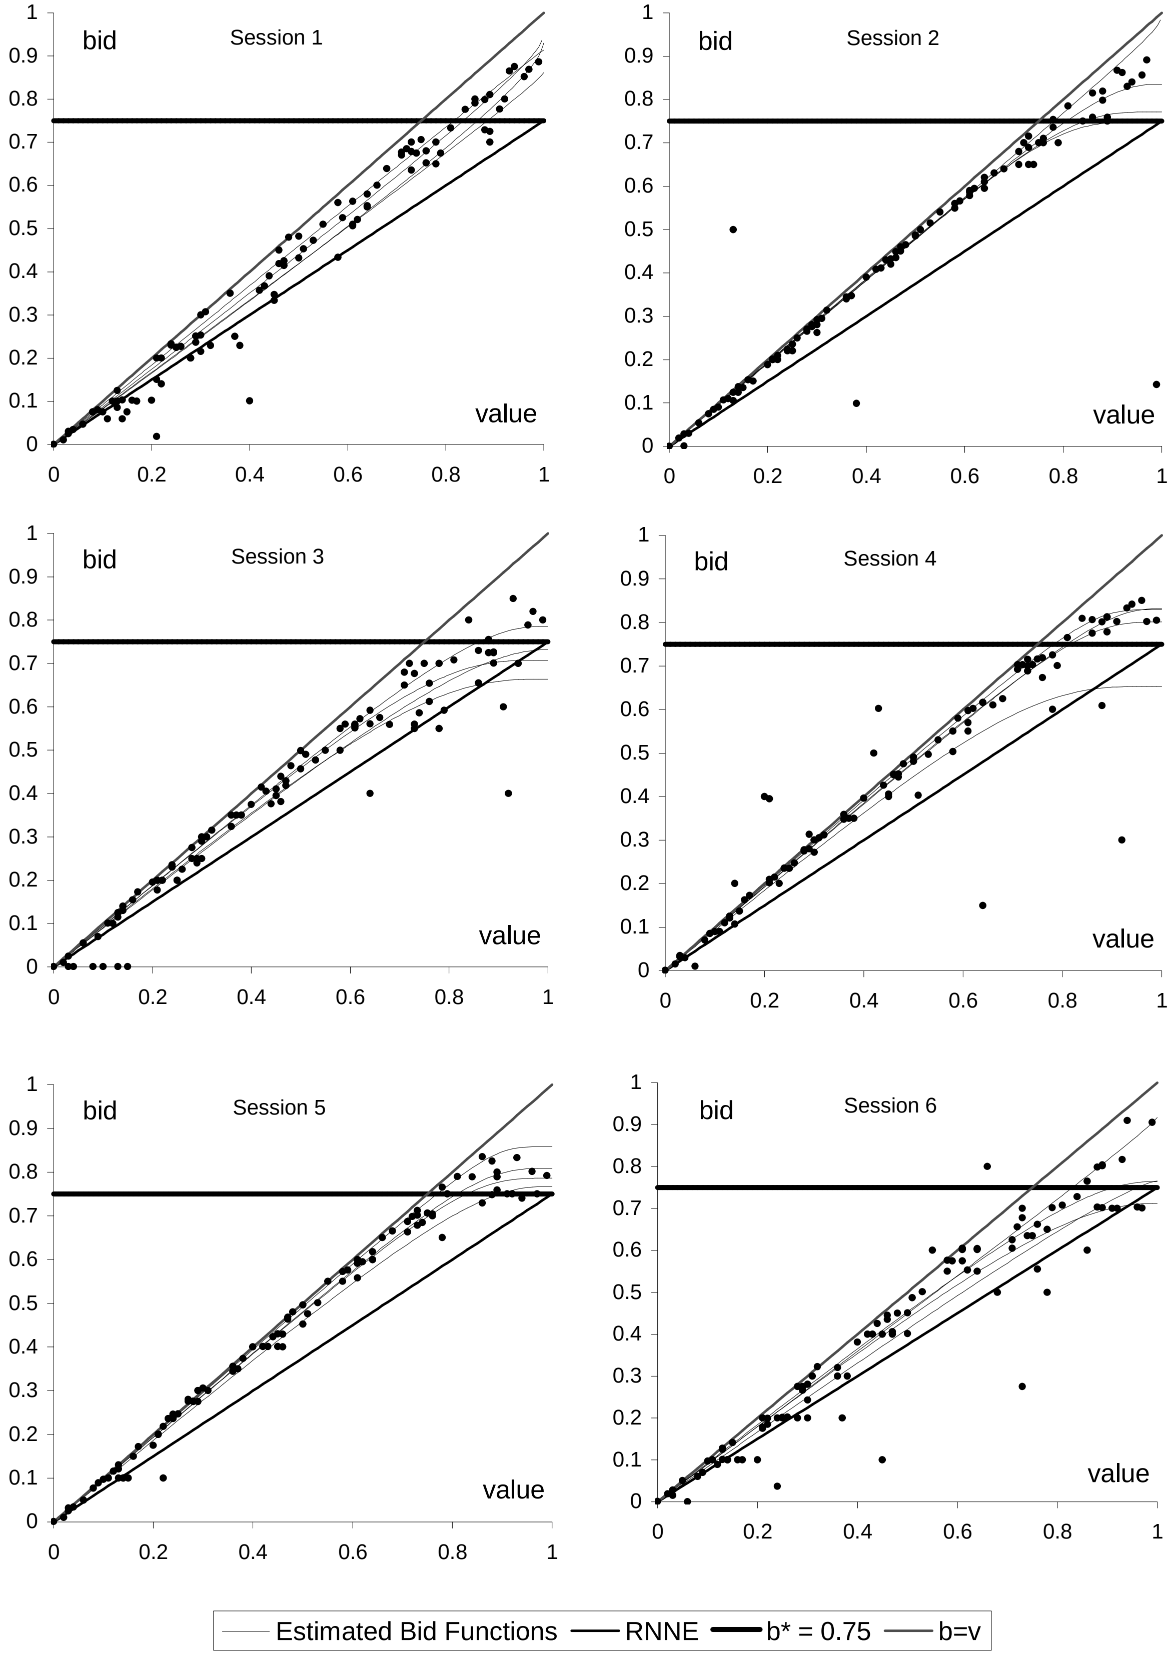

Figure 3: Estimated Individual Bid Functions

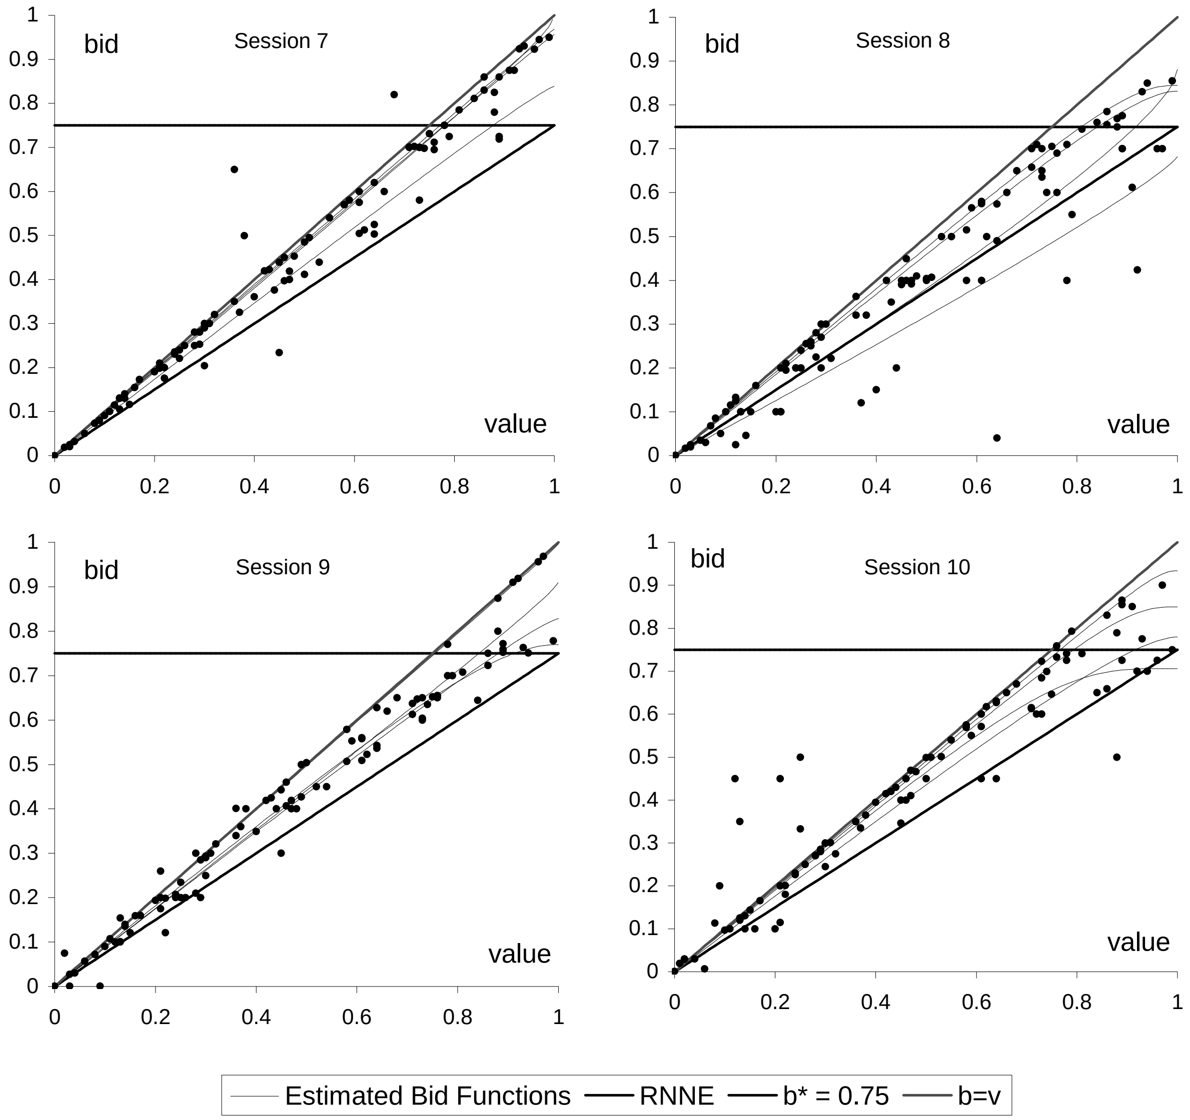

Appendix : Proof of Propositions

Proof of Proposition 1: Take $(\theta_1, 1, r)$, we will show that there exists another vector $\left(\theta_1^*, \theta_2^*, r^*\right)$ such that $\varphi(v;\theta_1^*, \theta_2^*, r^*) = \varphi(v;\theta_1, 1, 1)$ for all valuation *v*. Since $\theta_2 = 1$, we have that $F_\theta(s) = s^{\theta_1}$ and then,

$$\varphi(v;\theta_1, 1, r) = v - \int_0^v \left[\frac{s}{v}\right]^{\theta_1 \frac{n-1}{r}} ds =$$
$$= v - \frac{vr}{\theta_1(n-1)+r} = \varphi(v;\lambda\theta_1, 1, \lambda r)$$

for any λ. Therefore, there are infinitely many of these $\left(\theta_1^*, \theta_2^*, r^*\right)$. Moreover, note that if we take $\theta_1 = 1$ and $\lambda = 1/r$, we have that $\varphi(v;1,1,r) = \varphi(v;1/r,1,1)$ as put in claim 2. Although not needed in this paper, one further question is whether the result can be extended to the $\theta_2 = 1$ case, or in other words, whether the bidding function is one-to-one on the subspace of Θ where $\theta_2 = 1$ since in that case, local identification may arise**.** This subspace characterizes bidding functions which are not linear in valuations and a proof is not available. However, we have used numerical simulations to realize that it is always possible to find different vectors of parameters where the maximal distance between bidding functions is very small (less than 1e-20 in the worst case). These numbers would imply that bids functions are so close along these points that, even if there is local identification, the objective function of any estimation procedure would be too flat for an optimization algorithm to be numerically stable, which was the case when using the data in applications.

Proof of Claim 3: Non-identification arises because we are using first-order moments to estimate two parameters. Notice that m_W is the expectation of the second order statistic, i.e., $E(v_{(2)} | \theta_1, \theta_2)$. This expectation is taken with respect to an absolutely continuous distribution function. Thus, m_W is a continuous mappings of θ_1, θ_2 in their domain $\mathbb{R}^2_{++}$ and by continuity, for any θ_1, θ_2 there exists θ_1^*, θ_2^* such that $m_A(\theta_1, \theta_2, r) = m_A(\theta_1^*, \theta_2^*, r)$ and either $\theta_1 \neq \theta_1^{'}$ or $\theta_2 \neq \theta_2^{'}$. For m_A, the expectation of the bid function with respect to the Beta distribution, follows the same reasoning.

Proof of Proposition 2

To simplify notation, θ will denote the log-parameters which are denoted by $\tilde{\theta}$ in the text. Taking a Taylor expansion of $S_{nT}(\theta_0)$ around $\hat{\theta}$, we get

$$\frac{1}{\sqrt{nT}} S_{nT}(\theta^0) \# - \frac{1}{nT} \frac{\partial^2 S_{nT}(\theta^0)}{\partial^2 \theta} \sqrt{nT}(\hat{\theta} - \theta^0) \quad \text{(B1)}$$

By the central limit theorem, the left hand side of (B1) converges to *N(0,I₀)*. Further, since $\frac{1}{nT} \frac{\partial^2 S_{nT}(\theta^0)}{\partial^2 \theta} \xrightarrow{a.s.} J_0$ we have

$$\sqrt{nT}(\hat{\theta} - \theta^0) \xrightarrow{d} N(0, J_0^{-1} I_0 J_0^{-1}) \quad \text{(B1')}$$

Under the null hypothesis in (5), it holds that $h(r_0) = \theta^0$. Defining $S^*_{nT}(r) = S_{nT}(h(r))$, a Taylor expansion of $S^*_{nT}(r)$ around r_0 yields

$$0 = \frac{1}{\sqrt{nT}} \frac{\partial S^*_{nT}(\hat{r})}{\partial r} \# \frac{1}{\sqrt{nT}} \frac{\partial S^*_{nT}(r_0)}{\partial r} + \frac{1}{\sqrt{nT}} \frac{\partial^2 S^*_{nT}(r_0)}{\partial r \partial r'} (\hat{r} - r_0) \quad \text{(B2)}$$

Defining $h_{r_0} = \partial_r h(r_0)$, using $\frac{1}{nT} \frac{\partial S_{nT}(\theta^0)}{\partial \theta} \xrightarrow{a.s.} 0$ and re-arranging terms in (B2), we get

$$\sqrt{nT}(\hat{r} - r_0) \# (h_{r_0} J_0 h'_{r_0})^{-1} h_{r_0} \frac{1}{\sqrt{nT}} \partial_\theta S_{nT}(\theta^0) \quad \text{(B3)}$$

Using (B1), we have

$$\sqrt{nT}(\hat{r} - r_0) \# (h_{r_0} J_0 h'_{r_0})^{-1} h_{r_0} J_0 \sqrt{nT}(\hat{\theta} - \theta^0) \quad \text{(B4)}$$

Defining $\hat{\theta}^0 = h(\hat{r})$, a Taylor expansion of $\sqrt{nT}(h(\hat{r}) - h(r_0))$ around r_0 yields

$$\sqrt{nT}(\hat{\theta} - \theta^0) \# h_{r_0} \sqrt{nT}(\hat{r} - r_0) \quad \text{(B5)}$$

Hence, from (B5) and (B3), $\hat{\theta}^0$ and $\hat{\theta}$ are linearly related so that

$$\begin{aligned} \sqrt{nT}(\hat{\theta}^0 - \theta^0) &= h'_{r_0} (h_{r_0} J_0 h'_{r_0})^{-1} h_{r_0} J_0 \sqrt{nT}(\hat{\theta} - \theta^0) \\ &= P_h \sqrt{nT}(\hat{\theta} - \theta^0) \end{aligned} \quad \text{(B6)}$$

Since $\sqrt{nT}(\hat{\theta} - \hat{\theta}^0) = \sqrt{nT}(\hat{\theta} - \theta^0) - \sqrt{nT}(\hat{\theta}^0 - \theta^0)$, re-arranging terms in (B6) yields

$$\sqrt{nT}(\hat{\theta}-\hat{\theta}^0) = M_h\sqrt{nT}(\hat{\theta}-\theta^0) \quad \text{(B7)}$$

Now, θ^0, i.e. the true value of the parameters, is a vector (θ_1^0, θ_2^0), and under the null in (5) it must hold that $\theta_2^0 = 0$, while, $\theta_1^0 = -\ln(r_0)$. Thus, $h_{r_0}' = (-1/r_0, 0)$ and

$$M_h = \begin{pmatrix} 0 & \dfrac{-j_{12}}{j_{11}} \\ 0 & 1 \end{pmatrix} \quad \text{(B8)}$$

where j_{ik} stands for the element in row "*i*" and column "*k*" of J_0. From (B7) and (B8) it follows that

$$\sqrt{nT}\begin{pmatrix} \hat{\theta}_1 - \hat{\theta}_1^0 \\ \hat{\theta}_2 - \hat{\theta}_2^0 \end{pmatrix} = \sqrt{nT}\left(\hat{\theta}_2 - \theta_2^0\right)\begin{pmatrix} \dfrac{-j_{12}}{j_{11}} \\ 1 \end{pmatrix} \sim N\left(0, \sigma_{\hat{\theta}_2}^2 \begin{bmatrix} \left(\dfrac{-j_{12}}{j_{11}}\right)^2 & \dfrac{-j_{12}}{j_{11}} \\ \dfrac{-j_{12}}{j_{11}} & 1 \end{bmatrix}\right) \quad \text{(B9)}$$

Define

$$\aleph = nT(\hat{\theta}_1 - \hat{\theta}_1^0, \hat{\theta}_2 - \hat{\theta}_2^0)\sigma_{\hat{\theta}_2}^{-2} \begin{bmatrix} \left(\dfrac{-j_{12}}{j_{11}}\right)^2 & \dfrac{-j_{12}}{j_{11}} \\ \dfrac{-j_{12}}{j_{11}} & 1 \end{bmatrix}^+ \begin{pmatrix} \hat{\theta}_1 - \hat{\theta}_1^0 \\ \hat{\theta}_2 - \hat{\theta}_2^0 \end{pmatrix}$$

From (B9), the statistic $\aleph$ follows a chi-squared distribution with one degree of freedom (which is the rank of the variance matrix in (B9)) under the null. Using (B7) and (B8) and applying the definition of the Moore-Penrose generalized inverse, we eventually get

$$\aleph = \left[\frac{\hat{\theta}_2}{\sigma_{\hat{\theta}_2}}\right]^2$$

the square root of which follows a N(0,1) distribution. □