Playing like an Econometrician: Rational Choice Based on Estimates Leads to Overbidding in First-price Auctions

Leonardo Rezende*
PUC-Rio and University of Illinois

May 4, 2009

Abstract

Experiments with first-price auctions document that subjects tend to overbid in comparison to the risk-neutral Nash equilibrium. In this paper, I argue that overbidding may be done by procedurally rational agents that receive only partial information about behavior of other players and need to optimize over estimated, rather then true, bid distributions. I propose a *Best-Response to Estimate Equilibrium* and argue that, beyond overbidding, this concept captures a number of stylized facts observed in experimental data, such as the shape of the support, expectation, and variance of bids conditional on values.

1 Introduction

Studies involving experiments on first-price auctions, first carried out by Cox, Smith and Walker in the eighties, have found that subjects tend to systematically *overbid*, i.e., place bids that are significantly above the figures predicted by the symmetric Nash equilibrium when players are risk-neutral.

*Preliminary Version. Comments are most welcome. Address: Department of Economics, Catholic University at Rio de Janeiro (PUC-Rio). Rua Marquês de São Vicente, 225 SL 210F, Rio de Janeiro, Brazil, 22453-900. email: lrezende@econ.puc-rio.br. Tel.: 55-21-3527-1078, Fax: 55-12-3527-1084. All errors are mine.

This evidence poses a challenge to one of the fundamental models in auction theory, and have been addressed in a number of papers (Cox *et al.*, 1983, 1985, 1992; Goeree *et al.*, 2002; Harrison, 1989, 1990; Friedman, 1992; Kagel and Roth, 1992; Merlo and Schotter, 1992; Kagel and Levin, 1985; Armantier and Treich, 2005; Salo and Weber, 1995; Bajari and Hortaçsu, 2004; Rezende, 2006). Alternative explanations involve risk aversion (Cox *et al.*, 1985), weak incentives (Harrison, 1989), misperception of probabilities (Cox *et al.*, 1985; Goeree *et al.*, 2002; Armantier and Treich, 2005), "joy of winning" (Cox *et al.*, 1992; Goeree *et al.*, 2002) and ambiguity aversion (Salo and Weber, 1995).

In this paper, I propose a new explanation to the overbidding phenomenon. Here, I suggest that this behavior can be linked to *imperfect inference* about bidding by the other players.

In a Nash equilibrium one assumes that bidders are endowed with exact information about the other players strategies; in an incomplete information game such as an auction, one assumes that bidders know both the other players' type distribution but also the bid function, as well as knowing how to combine these pieces of information into a correct bid distribution. This seems to be a much more challenging cognitive task than to find the maximizing bid once that information is available. Additionally, and more importantly, even if bidders have infinite cognitive power, it is unclear how they would acquire the information needed about the other players behavior in equilibrium.

In this paper, I propose an alternative equilibrium concept that explicitly models imperfect inference about other players behavior. In a nutshell, I assume that each player must act like an econometrician before selecting her action; she observes a sample of bids drawn from the equilibrium bid distribution, forms estimates based on the sample, and places a bid that is the best response to beliefs based on those estimates. The equilibrium is a distribution that is a fixed point in this process, that is, when bids drawn from the equilibrium distribution are best response to estimates from samples from that same distribution. I call these *Best-Response to Estimates Equilibria*.

Since the inference adds some imprecision in the bidder's problem, clearly the actions chosen would deviate from the exact predictions from the Nash equilibrium. In this paper, I show that in the case of a First-price auction this deviation is systematic: players responding to estimates will tend to overbid compared to Nash *even when the estimation errors from the procedure they* *employ are not systematic!* This feature is in sharp contrast to to the literature based on probability misperception (Goeree *et al.*, 2002; Armantier and Treich, 2005), as there the overbidding is linked to a systematic pessimistic bias in the perceived probabilities (at least in the relevant range).

The papers is organized as follows: In section 2 I present the standard model of Nash behavior in a first-price auction and compare with experimental data from Guarino (2008). These data are typical of experiments with first-price auctions. I document systematic overbidding, and also point out a number of other empirical regularities yet to be explained.

In section 3 I present more formally the concept of best-response-to-estimate equilibrium. I also provide an existence result for this solution concept, based on the theory of Markov chains. In section 4 I apply the solution concept to the context of first-price auctions. I derive characterizations of best-response to estimate equilibria in the exact setting of the Guarino (2008) experiment, and use them to draw testable implications, which I test in section 5. Section 6 provides some concluding remarks.

2 Overbidding in First-Price Auctions

In this section, I review the standard theory of Nash equilibrium in first-price independent private value auctions and describe some empirical regularities found in data from an experiment involving these games.

In an symmetric independent private values first price auction $n+1$ players (called the bidders) compete to acquire a single item. The value of the item for bidder i is v_i, a random variable drawn from a given distribution F. The values for all bidders are drawn independently from the this same distribution. Bidder i observes his or her own value v_i, but not the other players values. Thus, this is an incomplete information game, where the values are the players' types.

The rules of the game asks for all players to simultaneously place bids. The bidder who has played the highest bid wins the item, and pays for it his or her own bid. Thus, writing b_i for i's bid, we find that player i earns a profit of $\pi_i = v_i - b_i$ if $b_i > b_j$, $\forall j \neq i$, and 0 otherwise. (If there is a tie, then a tie breaking rule is used. Typically, bidder i will earn $v_i - b_i$ with a probability equal to the inverse of the number of tying bidders, and 0 otherwise.)

Let $G_{NE}(t)$ be the cumulative distribution of a bid in a symmetric Nash equilibrium. To simplify notation, I assume this distribution is continuous, so that the probability of a tie is negligible. Knowing this distribution, bidder i selects a bid that maximizes expected profit:

$$(v_i - b_i)\Pr(b_i > b_j, \forall j \neq i) = (v_i - b_i)G_{NE}(b_i)^n.$$

Writing the solution to this problem as a function of v_i, one obtains the optimal bid function $\beta_{NE}(v_i)$. With knowledge of the bid function, one can compute the distribution $G_{NE}(t) = F(\beta_{NE}^{-1}(t))$. Observing that β_{NE} depends on the rivals' bid distribution, we find that the symmetric Nash equilibrium of this game can be characterized as a fixed point G_{NE} of the mapping that leads from the space of distributions into itself.

In what follows, I will be particularly interested in the case where the value distribution is uniform. In this case, the bid distribution arising in a Nash equilibrium will also be uniform. Suppose $b_j \sim U[a, c]$, the uniform distribution over the interval $[a, c]$. Then the best response for a bidder i with value v_i is

$$b_i = \begin{cases} \frac{n}{n+1}v_i + \frac{1}{n+1}a, & v_i \in [a, \frac{n+1}{n}c - \frac{1}{n}a]; \\ \text{any number in } [0, a], & v_i < a; \\ c, & v_i > \frac{n+1}{n}c - \frac{1}{n}a. \end{cases}$$

If $v_i \sim U[0, 1]$, I claim that in equilibrium $a = 0$, $c = \frac{n}{n+1}$. Indeed, for these choices always $v_i \in [a, \frac{n+1}{n}c - \frac{1}{n}a]$ and therefore $b_i = \frac{n}{n+1}v_i + \frac{1}{n+1}0 \sim U[0, \frac{n}{n+1}]$.

2.1 Experimental data

In this section, I present some statistical findings from experiments performed in Guarino (2008). I employ data from the first four sessions, which were the ones that involved subjects playing a first-price auction game exactly as described in the previous section. (The other sessions involved subjects playing against pre-specified automated strategies.)

In each session, 10 subjects were randomly grouped in pairs and each pair played an independent private values auction. (Thus, $n = 1$.) Values were drawn from the Uniform distribution between [0,10], and rounded up to the first decimal case. To facilitate comparison, I report data on values and bids normalizing to the unit interval. I also ignore the discretization and treat these values as continuous.

In each section 15 rounds of auctions were performed. In an effort to avoid repeated game effects, subjects were randomly re-matched after each repetition, and identities of opponents were kept secret. Overall, a sample of $4 \times 10 \times 15 = 600$ observations of value/bid pairs were obtained. In what follows I assume that repeated game effect are not present and that observations are mutually independent.

The experiments have been performed under the binary lottery procedure to control for risk attitudes (Roth and Malouf, 1979; Berg *et al.*, 1986; Cox *et al.*, 1985). In a nutshell, the idea is to pay bidders in lottery tickets instead of dollars so that behavior is not affected by risk aversion. If the winner of the auction receives x_1 with probability (proportional to) $v_i - b_i$ and x_2 otherwise, his or her payoff is $u(x_2) + [u(x_1) - u(x_2)](v_i - b_i)G(b_i)^n = A + B(v_i - b_i)G(b_i)^n$, and he or she should behave as a risk neutral bidder, independently of the shape of u. This is true as long as $B = u(x_1) - u(x_2) > 0$, that is, that x_1 is a better prize than x_2 for all individuals.

Figure 1 presents the data points with values in the horizontal axis and bids in the vertical axis. The flatter straight line is the locus $b = \frac{n}{n+1}v = v/2$ that is predicted in a Nash equilibrium; the steeper line is the 45° line. Clearly, bidders tend to overbid: 82.7% of the observed bids were above the Nash equilibrium prediction, while only 3% were equal to the latter. In comparison, 5% of the bids were equal to the value, and 1.2% were equal to 0, and 9.8% were equal to 5.

It is worth noting that, even though subjects typically overbid, they do sometimes underbid (in one extreme case, a bidder with value 0.89 places a bid of zero). Also, players never bid strictly above value. It is also apparent from the picture that the support of the empirical bid distribution is nearly as big as the one for values (the highest observed bid was 0.96).

Figure 2 presents plots of linear and quadratic regressions, as well as a non-parametric estimate of the conditional expectation of bids on values. The coefficients of the regressions are presented on table 1. Standard errors are presented in parentheses. The slope coefficient of the linear specification is 0.65, significantly above the Nash equilibrium prediction of 0.5. The intercept is positive, but not significant. The quadratic specification shows that, once again against the Nash equilibrium theory prediction, the linear functional form is rejected: The quadratic term is negative and significant.

Figure 2 presents a plot of a non-parametric estimate of the conditional expectation of bids on values, the Nadaraya-Watson local constant estimator.[1] The plot confirms the conclusion that the conditional expectation is

[1] This plot was done using a bandwidth of 0.1 and the Epanechnikov kernel.

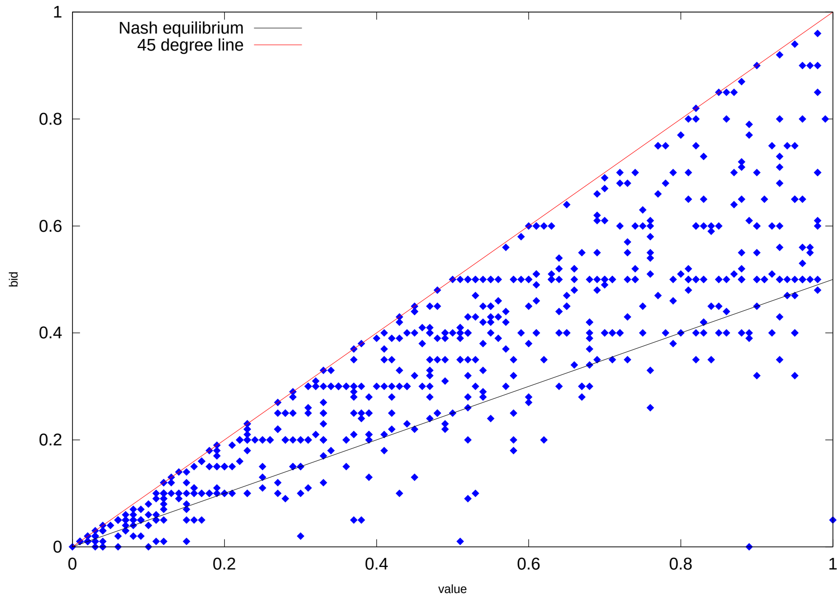

Figure 1: Guarino (2008) dataset, sessions 1–4.

intercept	v	v^2
0.01668	0.6565	
(0.00894)	(0.0155)	
-0.0207	0.8795	-0.2232
(0.0133)	(0.0610)	(0.0591)

Table 1: Regressions of bids on values.

concave. The relationship seems to be piecewise linear, with a kink at around 0.8.

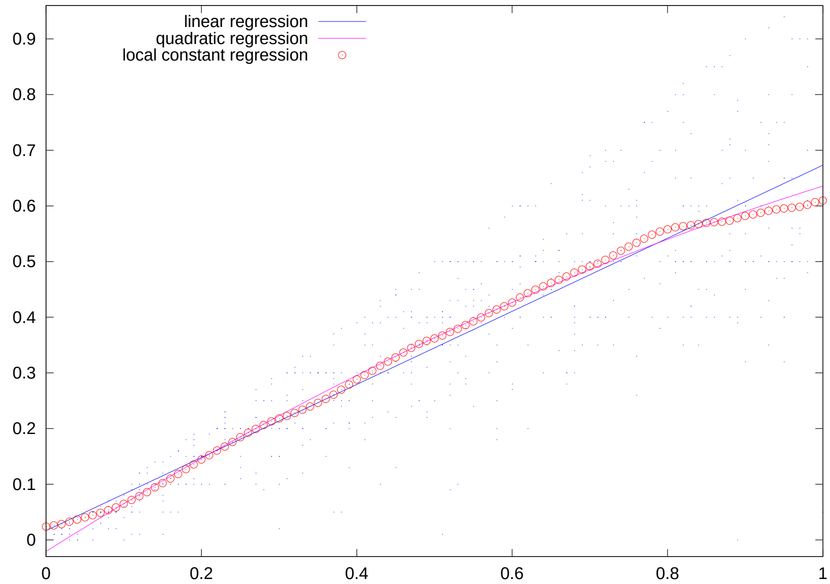

Figure 2: Estimates of the conditional expectation of bids on values.

Figure 1 also suggest a heteroskedasticy pattern in the data, with the conditional variance of bids being increasing in values. This is confirmed in figure 3, which shows a non-parametric estimate of the conditional variance of bids on values.[2] For comparison purposes, the figure also presents a plot of $v^2/12$, the conditional variance if $b|v \sim U[0, v]$.

We conclude this section by summarizing a number of stylized facts about the distribution of bids conditional on values:

1. On average, there is overbidding: most bids are above the Nash equilibrium prediction.

2. There is a positive probability of underbidding: even bidders with high values sometimes bid as low as 0.

[2]This is the Nadaraya-Watson estimator of the conditional expectation of squared residuals of the previous non-parametric regression, with the Epanechnikov kernel and a bandwidth of 0.2.

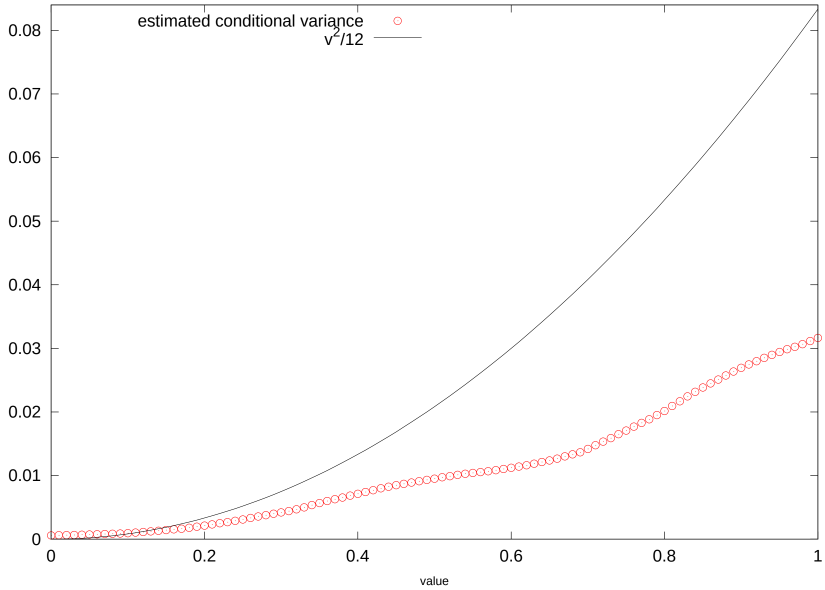

Figure 3: Estimate of the conditional variance of bids on values.

3. There is never bidding above value.
4. There is positive probability of bidding equal to value.
5. There is a positive probability of bids as high as the maximum value.
6. The conditional expectation of bids on values is concave, and everywhere above the Nash equilibrium prediction.
7. The conditional variance of bids on values is increasing.

3 Best-Response-to-Estimate Equilibrium

In this section, we provide a formal definition for the Best-Response-to-Estimate Equilibrium and discuss some of its general properties.

Let Γ be an incomplete information game with $n+1$ players. Let v_i represent the type of a player i and b_i her action in the game. We assume both v_i and b_i can be represented by real numbers. We also assume that Γ is a game with symmetric independent private types, namely, that $v_1, \ldots v_{n+1} \sim$ i.i.d.

F, and that the payoff of player i can be written as a function $\pi_i(b, v_i)$, where b is the profile of all actions. We also assume the game itself is symmetric, which means that the payoff function is the same across bidders, and the effect of actions of other players is exchangeable. These assumptions allow me to focus on symmetric strategy profiles.

It is standard practice to represent pure strategies as functions from v_i to b_i. Here, it is more convenient to characterize a strategy (pure or mixed) as the joint distribution $H(b_i, v_i)$ of action/type pairs. In a symmetric Nash equilibrium, all players select actions that generate the same distribution H_{NE}.

For a given joint distribution of action/types H, let $G(b_i)$ be the marginal distribution of actions associated with it. As before, let G_{NE} be the marginal distribution of actions in a Nash equilibrium.

For any distribution G over the space of action/type pairs, let[3]

$$\beta_i(G, v_i) = \underset{b_i'}{\operatorname{argmax}} \int \pi_i(b_i', b_{-i}, v_i) d \prod_{j \neq i} G(b_j).$$

This is the best response correspondence of a player i that believes her rivals will play actions that follow the distribution specified by G. In this notation the Nash Equilibrium distribution H_{NE} is such that, for every i and v_i, the distribution of b_i conditional on v_i according to H_{NE} equals the distribution of a selection of $\beta_i(G_{NE}, v_i)$.

I depart from the Nash equilibrium formulation by assuming that each player has access only to a sample of actions of the rivals.

Define a *sampling scheme* S to be a mapping between G and a distribution for a K-dimensional vector $X = (x_1, \ldots, x_K)$. X is a *sample* of G. An example of sampling scheme is random sampling, where X is a random vector i.i.d. G.

Define an *estimator* Ψ to be a mapping from a sample X to an estimate $\hat{G} = \Psi(X)$ for the distribution function for b.

Definition 1 (BE equilibrium) *A* Best-Response-to-an-Estimate Equilibrium *for a given sampling scheme S and choice of estimator Ψ is a distribution G such that, for each player i,*

[3] It is worth observing that, due to the private types assumption, H affects β_i only through the marginal distribution of actions; thus only information about G affect choice.

- *for each realization of S, i forms an estimate $\hat{G}$ according to the estimator Ψ;*
- *and there is a selection $\tilde{\beta}(\hat{G}, v_i) \in \beta(\hat{G}, v_i)$ that generates G; that is, $G(b_i) = \int \Pr(\tilde{\beta}(\hat{G}, v_i) \leq b_i | v_i) dF(v_i)$.*

The properties of a BE equilibrium naturally depend on the choice of sampling scheme and estimator being used. The next section will provide a characterization of the BE equilibrium to the case of a first-price auction. In the remainder of this section, I present a discussion of how existence proofs for this equilibrium concept can be obtained, as well as providing some remarks on the relation between BE equilibria and Nash Equilibria.

3.1 Existence

Existence of BE Equilibria can be shown exploring the theory of invariant distributions of Markov Chains.

One can associate a Markov Chain to the BE equilibrium construction as follows: Any sampling scheme leads to a mapping between a bid distribution G and the distribution of the sample X; for example, with random sampling this distribution is $\prod G$. I construct a Markov chain in the space of samples as follows: given a state X, let the next state X' to be a sample drawn from sampling scheme S on the bid distribution $\int \Pr(\beta(\psi(X))(v_i) \leq b_i | v_i) dF(v_i)$, generated by bids from players that best-respond to X.

By definition, a BE equilibrium is an invariant distribution of such Markov Chain. Any *recurrent* chain admits a unique invariant distribution (Meyn and Tweedie, 1993, Theorem 10.0.1). All that is required to prove existence and uniqueness of a symmetric BE equilibrium is to show that this chain is recurrent.

In the case where there are finitely many actions, recurrence corresponds to the condition that the probability of the chain going from any state to any other state eventually is positive; if P is the transition matrix for the chain, all elements of P^n should be positive for sufficiently large n. Even when the chain is not recurrent, it is often possible to restrict attention to an absorbing set (a set of states such that once in, the Markov chain never leaves) and verify recurrence in the absorbing set.

In the case of a continuum of actions, one must employ generalizations of the concept of recurrence. While recurrence is hard to verify directly, there a variety of sufficient conditions that can be easy to verify.

One example of set of sufficient conditions is as follows: Let $P(x, A)$ be the transition kernel of the Markov chain (the continuous space analog of the transition kernel), where x is a state and A a Borel set of future states. Write $L(x, A)$ for the probability that the chain ever enter set A starting from x. Following Meyn and Tweedie (1993), we call a Markov chain ψ-irreducible if there exists a measure ψ such that for every set A such that $\psi(A) > 0$, and every initial state x, $L(x, A) > 0$. In addition, we require a topological condition: A chain is (weak) Feller if, for any open set O, $P(\cdot, O)$ is a lower semicontinuous function of the initial state. Any ψ-irreducible Feller chain in a compact set is recurrent, and therefore has an invariant distribution.

3.2 Relation with Nash Equilibrium

The concept of BE equilibrium does not coincide with Nash equilibrium, except on the case where the bid distribution is degenerate. (Here by "Nash equilibrium" I refer to the joint distribution of actions and types induced by a symmetric Nash equilibrium, so that the two concepts are comparable.)

In general, the BE equilibrium will not be a Nash equilibrium, and vice-versa. Intuitively, sampling error will lead more dispersion on the actions in the BE equilibrium.

The concepts are related only for large sample sizes K. If the chosen estimators have good asymptotic properties, as $K \to \infty$ one could expect the BE equilibrium to approach the Nash equilibrium. However, the standard assumptions used to guarantee consistency of the estimator may not be enough, because here, unlike in the standard econometric problem, as the sample size K grows, so does the underlying distribution (i.e., the BE equilibrium associated with K).

4 BE equilibria in First-Price Auctions

In this section, I provide a characterization of a BE equilibrium in the case of a first-price auction with independent private values drawn from a uniform distribution.

In this setting, I suppose that a bidder i is required to select a bid b_i based on her own value v_i and a sample of equilibrium bids X. The first task is to select an estimation strategy. In this context, it is natural to work with the family of uniform distributions; after all, both the value distribution and the Nash equilibrium bid distribution belong to this family.

If the bidder believes that X is a random sample of the distribution $U[a, c]$, where a and c are unknown parameters, a natural choice of estimation procedure is maximum likelihood, which would lead to using $\underline{x} = \min X$ as an estimator for a and $\overline{x} = \max X$ as an estimator for c. From now, we assume that the estimation procedure is maximum likelihood based on the uniform distribution and characterize the BE equilibrium for this choice.

As discussed before, the best response for bids drawn from a uniform distribution is as follows:

$$\beta(\hat{G}, v_i) = \begin{cases} \frac{n}{n+1}v_i + \frac{1}{n+1}\underline{x}, & v_i \in [\underline{x}, \frac{n+1}{n}\overline{x} - \frac{1}{n}\underline{x}]; \\ \text{any number in } [0, \underline{x}], & v_i < \underline{x}; \\ \overline{x}, & v_i > \frac{n+1}{n}\overline{x} - \frac{1}{n}\underline{x}. \end{cases}$$

Before we proceed, two tasks remain; we must make a selection in the region where the best-response is not single-valued, and we must select a sampling scheme.

The best response is not single-valued if the bidder observes a sample where all bids are above the value. She infers that she cannot win the auction at any price below the value, and therefore cannot obtain positive profits in this game. Strictly speaking, any legal bid that leads to zero expected profits is optimal. Throughout the paper, I make the selection of a bid equal to the bidder's value. This is in line with equilibrium behavior of a bidder that believes she has the lowest possible value for the item, and makes the best-response function continuous on both v_i, $\underline{x}$ and $\overline{x}$. It is also compatible with the empirical evidence that a bid is never greater than, but is sometimes equal to, the value.

As for the sampling scheme, the obvious choice would be a random sample of k observations of the equilibrium bid distribution. However, this choice neglects an implicit additional source of information, the bidder's own value v_i. If i knows her own value v_i, and knows that the game is symmetric, she can infer that it is possible that all the other bidders have a value equal to her own. In a game where all bidders have values around v_i, equilibrium bids would also be around v_i. Thus, even in the absence of any other information, bidder i may infer that it is possible that some bid would be near v_i.

We provides two characterizations BE equilibria for two different sampling schemes. One sampling scheme is the "obvious choice" of a random sample $X = (x_1, \ldots, x_k)$, i.i.d. from the equilibrium bid distribution. The second scheme incorporates the implicit information discussed in the previous paragraph, using instead $X = (x_1, \ldots, x_k, v_i)$, where v_i is added to the estimation procedure as if it was a possible bid.

We call the BE equilibrium arising from the latter sampling scheme "BE equilibrium with introspection" (BEI) and the former "standard BE equilibrium" (BE). We characterize first the BEI, as it is analytically simpler.

4.0.1 BE equilibrium with introspection

In this section, we assume the bidder follows the estimation procedure outlined above, but utilizes a modified sampling scheme $X = (x_1, \ldots, x_k, v_i)$. Thus, the bidder always believes she could have seen a bid equal to her own value. We call that an "equilibrium with introspection" because a bidder could introspect that, if she has a value v_i, all other bidders could have that value as well, in which case one would observe an auction price equal to v_i.

With this modified sampling scheme, it is always true that $v_i \in [\underline{x}, \overline{x}]$. Thus, the optimal bid simplifies to

$$b_i = \frac{n}{n+1} v_i + \frac{1}{n+1} \underline{x}.$$

We can immediately observe that, for any realization of v_i, $b_i \geq \beta_{NE}(v_i) = \frac{n}{n+1} v_i$. Also, $b_i \leq v_i$. We can anticipate that the bid distribution of the modified BE equilibrium, G_{BEI}, will first-order stochastically dominate G_{NE}, and will be dominated by F.

To characterize G_{BEI}, we first develop an expression for the distribution of $\underline{x}$ conditional on v_i, as it affects the bid. To save notation, we drop the i subscript of v_i and b_i. We have, if $t < v$,

$$\Pr(\underline{x} > t|v) = \Pr(x_1 > t|v) \times \cdots \times \Pr(x_k > t|v) = (1 - G_{BEI}(t))^k;$$

if $t \geq v$, $\Pr(\underline{x} > t|v) = 0$; and of course, if $t \leq 0$, $\Pr(\underline{x} > t|v) = 1$. From this, we can compute the distribution of b conditional on v:

$$\begin{aligned} \Pr(b \leq t|v) &= \Pr(\underline{x} \leq (n+1)t - nv) \\ &= \begin{cases} 0 & , t \leq \frac{n}{n+1} v \\ 1 - (1 - G_{BEI}((n+1)t - nv))^k & , \frac{n}{n+1} v < t < v \\ 1 & , t \geq v \end{cases} \end{aligned}$$

Integrating over the distribution of v, we find the unconditional distribution of b, which in equilibrium must coincide with G_{BEI}. If $t < \frac{n}{n+1}$,

$$\begin{aligned}
G_{BEI}(t) &= \int_0^1 \Pr(b \le t|v)dv \\
&= \int_0^t 1\, dv + \int_t^{\frac{n+1}{n}t} [1-(1-G_{BEI}((n+1)t-nv))^k]dv \\
&= \tfrac{n+1}{n}t - \textstyle\int_t^{\frac{n+1}{n}t}(1-G_{BEI}((n+1)t-nv))^k dv \\
&= 1 - \frac{1}{n}\int_0^t (1-G_{BEI}(s))^k ds;
\end{aligned}$$

if $t > \frac{n}{n+1}$,

$$\begin{aligned}
G_{BEI}(t) &= \int_0^1 \Pr(b \le t|v)dv \\
&= \int_0^t 1\, dv + \int_t^1 [1-(1-G_{BEI}((n+1)t-nv))^k]dv \\
&= 1 - \int_t^1 (1-G_{BEI}((n+1)t-nv))^k dv \\
&= 1 - \frac{1}{n}\int_{(n+1)t-n}^t (1-G_{BEI}(s))^k ds
\end{aligned}$$

We have therefore an integral equation

$$G_{BEI}(t) = \begin{cases} \frac{n+1}{n}t - \frac{1}{n}\int_0^t (1-G_{BEI}(s))^k ds & , t < \frac{n}{n+1} \\ 1 - \frac{1}{n}\int_{(n+1)t-n}^t (1-G_{BEI}(s))^k ds & , t \ge \frac{n}{n+1} \end{cases}$$

that along with the initial condition $G_{BEI}(0) = 0$ characterizes G_{BEI}. At $t \in [0, \frac{n}{n+1}]$ this is equivalent to the nonlinear autonomous ordinary differential equation $G'_{BEI}(t) = \frac{n+1}{n} - \frac{1}{n}(1-G_{BEI}(t))^k$. In the region $t \in [\frac{n}{n+1}, 1]$, this becomes a delay differential equation.

From the fact that $G_{BEI}(s) \in [0,1]$, using the integral equation one can readily verify that $t \le G_{BEI}(t) \le G_{NE}(t)$, where G_{NE} is the Nash equilibrium bid distribution.

4.0.2 Standard BE equilibrium

We now consider the case where introspection does not occur and the bidder forms her estimate using a sample X including only k independent bids drawn from G_{BE}.

In this case, since the optimal bid depends on both $\underline{x}$ and $\overline{x}$, we need to employ the joint distribution of these variables, which are now independent of v.

We do so by observing that $\Pr(\overline{x} \leq s) = G_{BE}(s)^k$, and that, for $t < s$,

$$\Pr(\underline{x} \leq t | \overline{x} \leq s) = 1 - (Pr(x_j \in (t, s] | x_j \leq s)^k = 1 - \left(\frac{G_{BE}(s) - G_{BE}(t)}{G_{BE}(s)} \right)^k,$$

and for $t \geq s$, $\Pr(\underline{x} \leq t | \Pr(\overline{x} \leq s)) = 1$. Combining the two expressions we obtain

$$\Pr(\underline{x} \leq t, \overline{x} \leq s) = \begin{cases} G_{BE}(s)^k - (G_{BE}(s) - G_{BE}(t))^k, & t < s; \\ G_{BE}(s)^k, & t \geq s. \end{cases}$$

Similarly, $\Pr(\underline{x} \geq t, \overline{x} \leq s) = (G_{BE}(s) - G_{BE}(t))^k$, if $s > t$, and 0 otherwise.

Using these expressions, we obtain the following distribution for a bid, conditional on the value v:

$$\Pr(b \leq t | v) = \begin{cases} G_{BE}(t)^k & , t \leq \{ \frac{n}{n+1} v \\ 1 - (1 - G_{BE}((n+1)t - nv))^k & \\ \quad + (G_{BE}(t) - G_{BE}((n+1)t - nv))^k & , \frac{n}{n+1} v < t \leq v \\ 1 & , t > v \end{cases}$$

This distribution is discontinuous at $t = v$; $\Pr(b = v | v) = (1 - G_{BE}(v))^k$. This atom is due to the assumption that $b = v$ when $\underline{x} > v$. Elsewhere (including at $\frac{n}{n+1} v$) it is continuous (assuming of course g_{BE} is continuous).

To obtain an integral equation for G_{BE}, we integrate $\Pr(b \leq t | v)$ over the distribution of v. If $t < \frac{n}{n+1}$,

$$\begin{aligned} G_{BE}(t) &= \tfrac{n+1}{n} t + (1 - \tfrac{n+1}{n} t) G_{BE}(t)^k \\ &\quad - \frac{1}{n} \int_0^t (1 - G_{BE}(s))^k - (G_{BE}(t) - G_{BE}(s))^k dv; \end{aligned}$$

and if $t \geq \frac{n}{n+1}$,

$$G_{BE}(t) = 1 - \frac{1}{n} \int_{(n+1)t-n}^t (1 - G_{BE}(s))^k - (G_{BE}(t) - G_{BE}(s))^k dv.$$

5 Do People Play like Econometricians?

In this section, we report result of simulated BE equilibria and investigate if they produce predicted statistics that are similar to the ones found in the experimental data. To assess the sensitivity of the results to the size of the sample observed by the bidders, results for two values for k are reported: a small value of $k = 3$, and a moderate value of $k = 10$. (For large values of k, the BE equilibrium bids are approximately equal to the Nash equilibrium.)

Figures 4 and 5 present plots of the predicted bid distributions for the BE equilibrium (labeled BE) and the BE equilibrium with introspection (BEI). The steeper straight line is the bid distribution in the Nash equilibrium and the flatter straight line is the value distribution. The picture also presents the empirical bid distribution in the Guarino dataset.

The BEI distribution always first-order stochastically dominates the Nash equilibrium distribution. The same is not true for the BE distribution; for every value of k, G_{BE} is above G_{NE} in a neighborhood below 0.5, although the neighborhood becomes small as k grows. Notwithstanding this, the expected bid under the BE equilibrium is above the Nash equilibrium in all cases except when $k = 2$, as reported in table 2.

Still, the amounts of predicted overbidding are smaller than in the data (the average bid is 0.34495). Figures 4 and 5 show that predicted distributions are similar in shape to the empirical one, but tend to lie to the left of it. Also, in the data bidders with low values tend to bid less aggressively, and bidder with high values more aggressively than predicted by the theory.

The theory also predicts some specific patterns in the joint distribution of bids and values that are qualitatively similar to those in the data. Both the BE and BEI predict that bids always lie at or below values, and this is verified in the data. Less obviously, both the BE and BEI predict that there is a strictly positive probability of observing a bid *equal* to value, as there is an atom at this point in the conditional distribution of bid given value. This apparently irrational choice is indeed observed in the data. Furthermore, theory predicts that the event $b = v$ is more likely when v is small; this pattern is verified in the data as well. Figures 6 and 7 plot the predicted and estimated densities[4] conditional on this event and both are generally decreasing.

Figures 8 and 9 present conditional expectations of bids given values.

[4]For the estimated density a Gaussian kernel with bandwidth 0.05 was used.

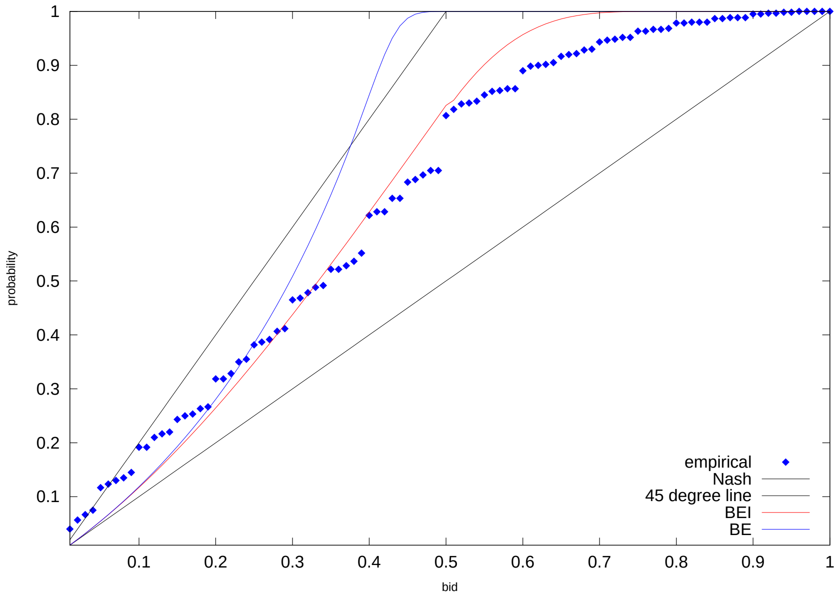

Figure 4: Estimated and Predicted Bid distributions, $k = 3$.

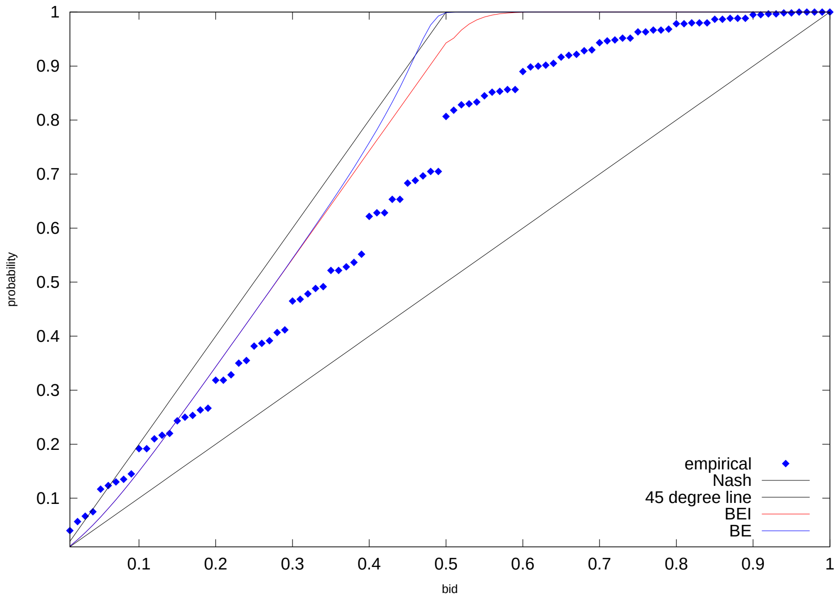

Figure 5: Estimated and Predicted Bid distributions, $k = 10$.

k	NE	BEI	BE
2	0.25	0.34613	0.23450
3	0.25	0.32191	0.26818
4	0.25	0.30690	0.27417
5	0.25	0.29663	0.27414
6	0.25	0.28916	0.27265
7	0.25	0.28348	0.27068
8	0.25	0.27902	0.26885
9	0.25	0.27541	0.26717
10	0.25	0.27244	0.26556
11	0.25	0.26995	0.26420
12	0.25	0.26784	0.26290
13	0.25	0.26602	0.26173
14	0.25	0.26443	0.26075
15	0.25	0.26305	0.25979
16	0.25	0.26182	0.25891
17	0.25	0.26072	0.25812
18	0.25	0.25974	0.25746
19	0.25	0.25886	0.25679
20	0.25	0.25806	0.25618

Table 2: Predicted Expected bids.

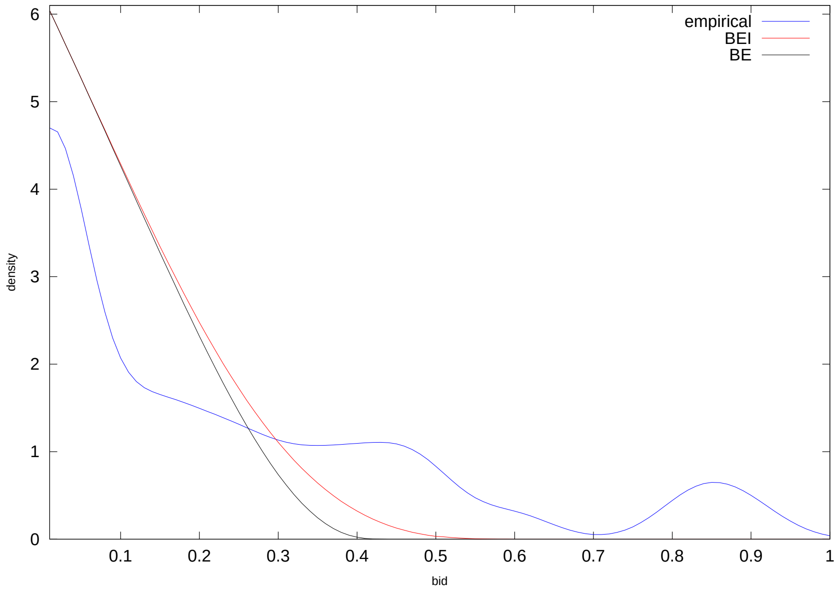

Figure 6: Bid density conditional on the event $b = v$, $k = 3$.

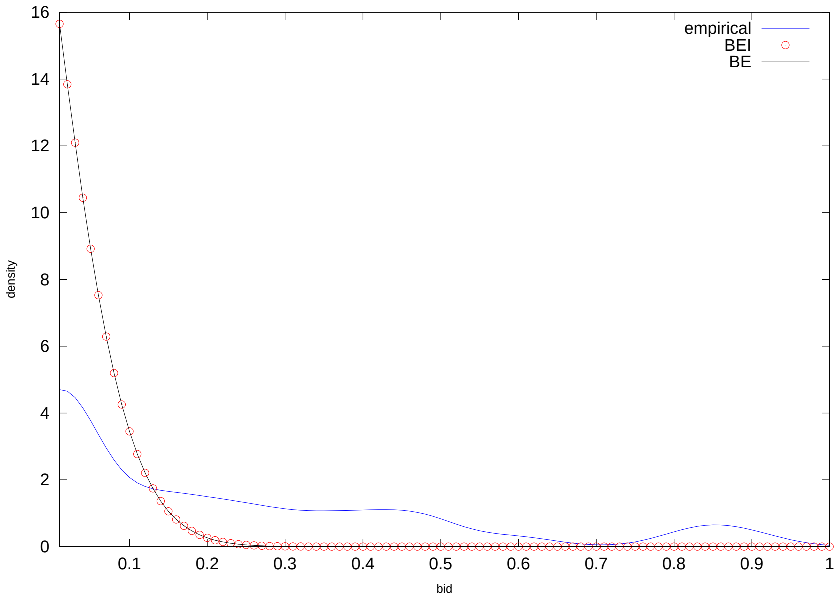

Figure 7: Bid density conditional on the event $b = v$, $k = 10$.

In line with the findings concerning the marginal bid distributions, we find that there is generally overbidding, but less so than in the data. In the BEI equilibrium, bidders with all values are expected to overbid. In the BE equilibrium, all bidders except those with highest values tend to overbid. Compared to the data, we see that again the theory predicts higher expected bids for subject with very low values.

Both the BEI and the BE equilibrium predicted expected valuation that are concave, as found in the data. Remarkably, for moderate values of k the BE equilibrium predicts a near kink that resembles the one observed in the data, as can be seen in figure 9.

The documented pattern of increasing heteroskedasticity is also predicted by the theory of best-response to estimate equilibrium. Figures 10 and 11 present the predicted conditional variance of bids given values, and compares them with the estimates ones. All of them are increasing, although predicted variances are much smaller in the case of the BEI equilibrium and are larger in the case of the BE equilibrium.

The BEI equilibrium predicts that bids are never below the Nash equilibrium bids (and never above values); since subjects do bid outside this range, the dispersion tends to be larger in the data than in theory. On the other hand, in the BE equilibrium, even bidders with high values place very low bids with some probability. This is indeed observed in the data, but less frequently then predicted by theory.

6 Concluding Remarks

In summary, one can conclude from the discussion in section 5 that the theory proposed in this paper is successful in predicting *qualitative* features in the observed patterns of behavior. It predicts overbidding; it predicts that subject never bid above value, but sometimes bid at value, and sometimes underbid as low as zero; it predicts that the conditional expectation of bids given values is increasing and concave; it predicts that the conditional variance is increasing; and it predicts that bids equal to values are more likely to appear when values are low. All of these features are found in the data.

Quantitative predictions are less successful; for example, the predicted intensity of overbidding is still smaller than what is found in that data. One should keep in mind however that the results reported in this paper are entirely out of sample; no information from the sample has been used to select

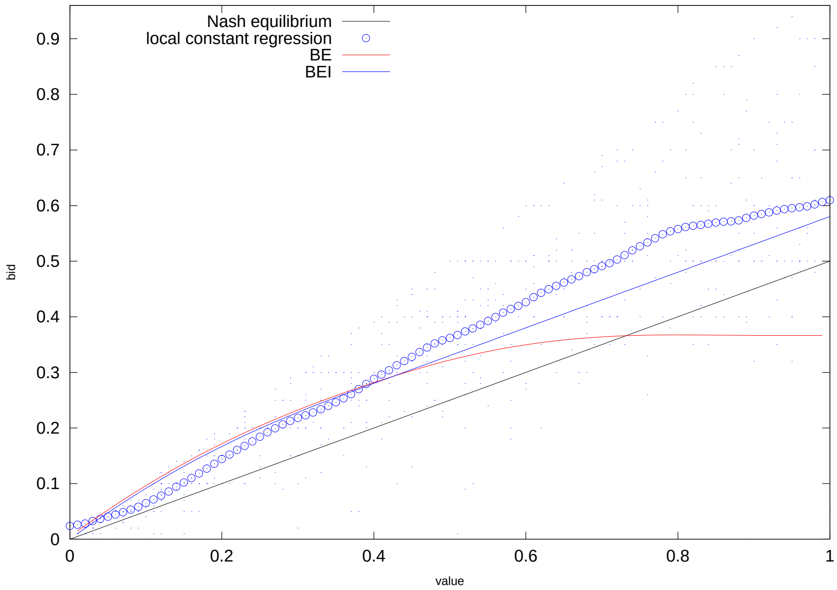

Figure 8: Conditional expectations of bids given values, $k = 3$.

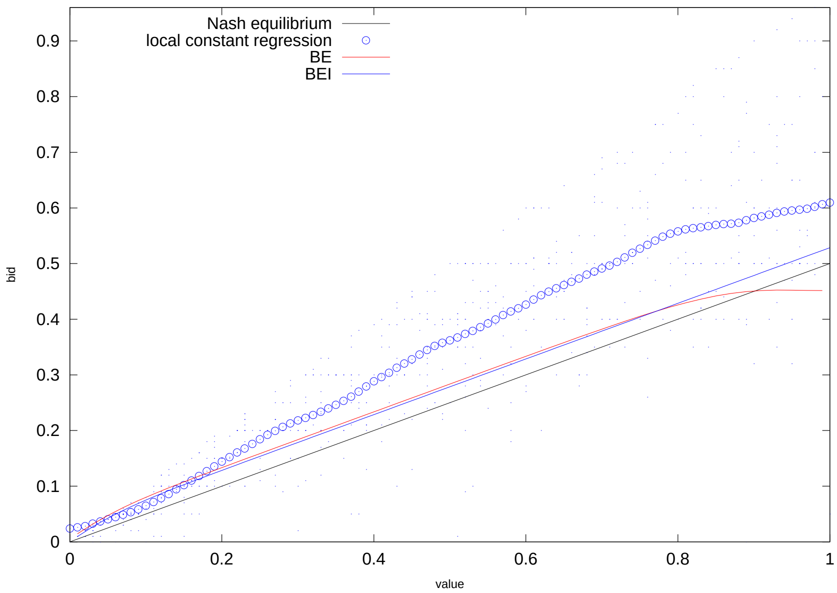

Figure 9: Conditional expectations of bids given values, $k = 10$.

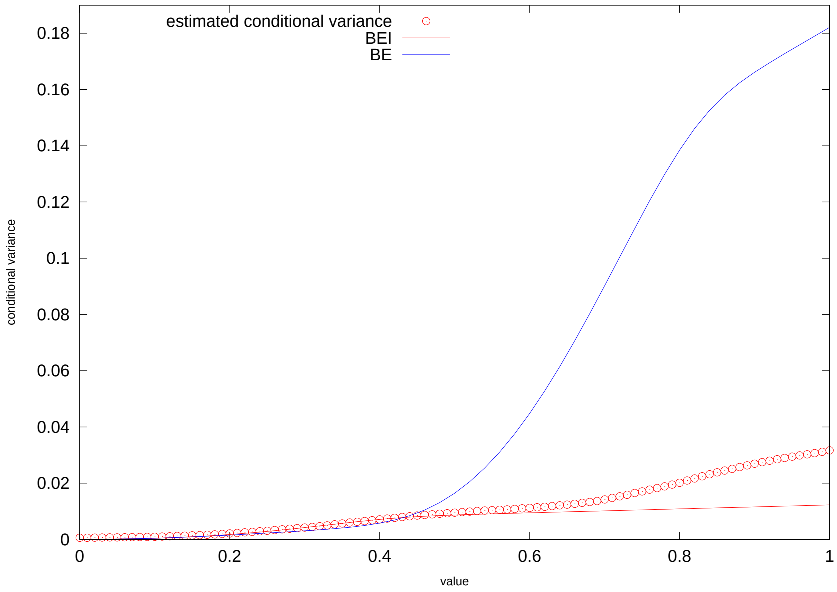

Figure 10: Conditional variances of bids given values, $k = 3$.

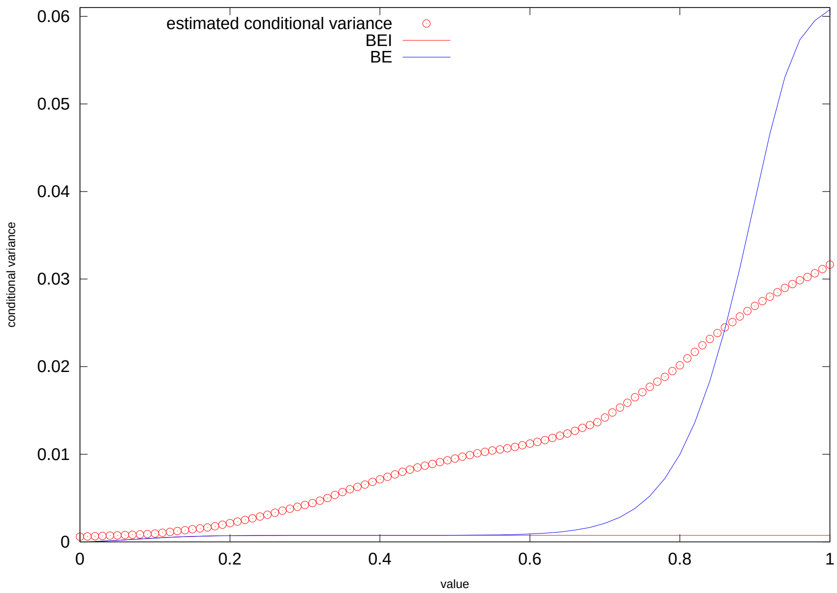

Figure 11: Conditional variances of bids given values, $k = 10$.

or calibrate the BE equilibrium predictions. It would indeed be remarkable that the somewhat arbitrary choices made in section 4 would lead to optimal goodness of fit.

Perhaps the task of obtaining a quantitatively accurate description of bidder behavior in this auction may be accomplished by searching over BE equilibria with variety of sampling and estimation procedures; perhaps other ingredients are needed, such as risk aversion or probability misperception.

However, this is not the main task of the paper; rather, the objective here was to propose a framework to formalize procedurally rational behavior and to verify that it can provide explanations from some apparent departures from rationality in one experimental setting. A more fruitful avenue of future research is to investigate if the Best-response to Estimate equilibrium framework can be successful in explaining anomalies found in experiments about other games as well.

References

Armantier O, Treich N (2005). Overbidding in Independent Private Values Auctions and Misperception of Probabilities. Technical report, Mimeo, Universite de Montreal.

Bajari P, Hortaçsu A (2004). Are structural estimates of auction models reasonable: Evidence from experimental data, duke and Chicago. Forthcoming on the *Journal Of Political Economy*.

Berg JE, Daley LA, Dickhaut JW, O'Brien JR (1986). Controlling preferences for lotteries on units of experimental exchange. *Quarterly Journal of Economics* **101**:281–306.

Cox JC, Smith VL, Walker JM (1983). Tests of a heterogeneous bidders theory of first price auctions. *Economics Letters* **12**:207–212.

Cox JC, Smith VL, Walker JM (1985). Experimental development of sealed-bid auction theory; calibrating controls for risk aversion. *American Economic Review Papers and Proceedings* **75**:160–165.

Cox JC, Smith VL, Walker JM (1992). Theory and misbehavior of first-price auctions: Comment. *American Economic Review* **82**:1392–1412.

Friedman D (1992). Theory and misbehavior of first-price auctions: Comment. *American Economic Review* **82**:1374–1378.

Goeree JK, Holt CA, Palfrey TR (2002). Quantal response equilibrium and overbidding in private-value auctions. *Journal of Economic Theory* **104**:247–272.

Guarino I (2008). *Identificando as Causas do Fenômeno de Overbidding em Leilões de Primeiro Preço: Uma Análise Experimental (Identifying the Sources of Overbidding in First-Price Auctions: An Experimental Analysis).* Master's thesis, Pontifical Catholic University of Rio de Janeiro (PUC-Rio).

Harrison GW (1989). Theory and misbehavior of first-price auctions. *American Economic Review* **79**:749–762.

Harrison GW (1990). Risk attitudes in first-price auction experiments: A bayesian analysis. *Review of Economics and Statistics* **72**:541–546.

Kagel JH, Levin D (1985). Individual bidder behavior in first-price value auctions. *Economics Letters* **19**:125–128.

Kagel JH, Roth AE (1992). Theory and misbehavior of first-price auctions: Comment. *American Economic Review* **82**:1379–1391.

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

Meyn S, Tweedie R (1993). *Markov chains and stochastic stability.* London: Springer-Verlag.

Rezende L (2006). Identifying the source of misbehavior in first-price auctions with bounds on overbidding, University of Illinois.

Roth AE, Malouf MWK (1979). Game-theoretic models and the role of information in bargaining. *Psychological Review* **86**.

Salo A, Weber M (1995). Ambiguity aversion in first-price sealed-bid auctions. *Journal of Risk and Uncertainty* **11**:123–137.