DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

QUANTAL RESPONSE AND NONEQUILIBRIUM BELIEFS EXPLAIN OVERBIDDING IN MAXIMUM-VALUE AUCTIONS

Colin Camerer

Salvatore Nunnari

Thomas R. Palfrey

SOCIAL SCIENCE WORKING PAPER 1349

June 2011

Quantal Response and Nonequilibrium Beliefs Explain Overbidding in Maximum-Value Auctions*

Colin Camerer   Salvatore Nunnari   Thomas R. Palfrey
California Institute of Technology

January 13, 2011

Abstract

We analyze data from the second-price maximum-value auction experiment reported in Ivanov, Levin and Niederle (2010) in order to investigate the extent to which the winner's curse can be explained by quantal response, in combination with diffferent assumptions about equilibrium (QRE) or nonequilibrium beliefs (QCH or QCE). We find a close correspondence between the theoretical predictions of those models and experimental behavior, even in the presence of frequent (and extreme) overbidding. The basic pattern in the data consists of a combination of flat overbidding for low signals, and monotonically increasing bidding for higher signals. The logit QRE model fits this pattern reasonably well. Incorporating nonequilibrium beliefs into the QRE model, in the form of either levels of strategic sophistication or cursed beliefs, leads to an even better match with the data, as these models imply slightly higher bids. We also show that these imperfect best response models predict essentially no treatment effects across different versions of the game, consistent with the experimental findings. Overall, our study indicates that the winner curse phenomenon in this auction is plausibly attributable to limits on strategic thinking combined with imperfect best response.

*We thank Asen Ivanov for promptly supplying the original data from the experiments in Ivanov, Levin and Niederle (2010), and to audience members of ESA Meeting Fall 2010, especially Muriel Niederle, for comments. Please do not quote or circulate without permission.

1 Introduction

This paper discusses experimental evidence of bidding in a maximum-value auction game. In these game, two bidders draw uniformly-distributed, independent, integer-valued private signals from the set $\{0, 1, \ldots, 10\}$. The common value of the auctioned object is the maximum of the two signals. The two bidders make simultaneous bids in a second-price auction. The highest bidder wins the object and pays the second-highest price. Ties are broken randomly.

The maximum-value auction has some interesting properties for testing theories of behavior in games. There is a unique symmetric Bayesian-Nash equilibrium in which bidders bid their signals. However, the equilibrium is weak because overbidding against signal-bidding is also a best response. Therefore, regular quantal response equilibria (Goeree, Holt and Palfrey (2005)) will typically entail a substantial frequency of overbidding.

Ivanov, Levin and Niederle (2010b) (ILN) conducted experiments on this maximum-value auction game. They also conducted two different experimental treatments: in one treatment (*MinBid*) bids cannot be lower than private signals; in the other treatment (*phase II*), subjects bid against their own previous bids from a *phase I* auction against other people.[1]

These are the two main results: first, there is substantial overbidding; the median bids in phase I of the experiment range monotonically from about \$5 to \$10 for corresponding signal values from \$0 to \$10. Second, bids do not change very much between the phase I auctions and the MinBid and phase II treatments.

We show that these two empirical results are consistent with both quantal response equilibrium (QRE) and with two theories in which strategic thinking is limited, cursed equilibrium (CE) and cognitive hierarchy (CH, similar to level-k), when quantal response is also assumed.

These results do not arise because 'quantal response can fit anything', a claim that is not generally true for regular quantal response equilibria including the most common logit and power specifications (Goeree, Holt and Palfrey, 2005).[2] Instead, the phase I data are first used to derive parameters which fit those data best using maximum likelihood estimation. Then the phase II bidding results are forecasted using those same parameters derived from

[1] There are also two versions of phase II, reminding subjects of their phase I bids (*ShowBid*) or not reminding them (*Baseline*). We pool those two versions because the results are similar.

[2] Regular equilibria satisfy various properties of continuity and responsiveness of choice probabilities to payoffs, and monotonicity– higher expected payoffs lead to higher choice probabilities.

fitting phase I.[3] Phase II bids forecasted out-of-sample and phase I bids fitted in-sample are predicted to be very similar (actually, predicted bids are slightly lower in phase II)– that is, the treatment effects predicted based on the phase I behavior are very small. And in fact, the empirical bid functions are also very similar. We also compute out of sample forecasts for the *MinBid* treatment, with similar results. Thus, the pattern of bids in the initial phase I auctions and the empirical (non-)response to the *MinBid* and phase II auctions can be empirically explained by QRE, and can be explained somewhat more accurately by quantal response versions of the CE and CH models.

1.1 The Winner's Curse

The impetus for ILN's paper is to further understand the winner's curse (WC). The winner's curse results when bidders bid too high a fraction of a private value estimate, in common-value auctions, leading them to overpay (relative to equilibrium bidding) and even lose money. The winner's curse is one of the most interesting and best-documented facts in empirical auction analysis (Kagel and Levin (1986), Dyer, Kagel and Levin (1989), Levin, Kagel and Richard (1996), Kagel and Levin (2002)). The fundamental explanation for the observed overbidding has, surprisingly, proved to be elusive.

A likely cause contributing to overbidding is the failure of agents to appreciate the relation between the bids of other agents and those agents' perceived value of the auctioned good. This mistake is the motivation for the concept of "cursed equilibrium" (Eyster and Rabin (2005); CE) and similar limits on rationality which motivate analogy-based expectation equilibrium (Jehiel (2005), Jehiel and Koessler (2008)).

Another possible cause is that agents simply find it conceptually difficult to compute a conditional expectation of value of any sort. If the difficulty of this computation is the main source of the winner's curse, the same mistake should also occur in a single-person decision where there is no other player or strategic uncertainty, but there is identical conditional expectation. Indeed, an early control experiment to the "Acquire-a-company" game conducted by psychologists showed that failures to adjust for adverse selection do also occur reliably in

[3] Note that if the models were wrong, and were overfit to phase I data, then the phase II and MinBid predictions would be farther off than the phase I fits, and they are not. Another common procedure is to fit models in each phase using part of the data, then predict for a hold-out validation sample. We have not done this but conjecture that the results would be similar, except for the caveat that obvious heterogeneity in the data implies that hold-out forecasts for some outlying subjects will be poor.

an individual-decision isomorph of that game (the "beastie run"; see Carroll, Bazerman and Maury (1988)). This fact became well known in decision research in the late 1980s and was frequently discussed in that literature (see recently Tor and Bazerman (2003), who show the relationship to imperfect updating in the Monty Hall 3-door decision problem). Charness and Levin (2009) report evidence from a similar decision experiment, replicating these earlier findings in psychology that winner's curse-like effects are also present in individual decision problems.

What about strategic models of the winner's curse? Besides CE, a class of non-equilibrium models, introduced in behavioral game theory around 1995, which have been applied to many types of games, allow players' beliefs about the likely actions of others to be mistaken. One such model, mentioned above, is cursed equilibrium. Another class of models include "level-k" and "cognitive hierarchy" (CH) specifications of beliefs.[4] Both models start with some clearly specified level-0 play (which is typically assumed to be random and uniform, but could more generally be considered instinctive (Rubinstein, 2007), focal based on salience (Crawford and Irriberri, 2007; Milosavljevic, Smith, Koch and Camerer, 2010).

Level-1 players assume they are playing level-0 players.[5] An iterated hierarchy then follows, most typically in two different ways: in "level-k" models level-k players (for $k \geq 1$) believe their opponents are all using level $k-1$ reasoning. In cognitive hierarchy (CH) models level-k players believe opponents are using a mixture of level $0, 1, \ldots, k-1$ reasoning; their normalized beliefs therefore approach equilibrium beliefs as k increases.

The simplest forms of these models assume players choose best responses given their beliefs. Indeed, ILN use the term "belief-based models" to refer exclusively to models with nonequilibrium beliefs and best response. They write:

> "[previous papers] rationalize the WC [winner's curse] within theories that retain the BNE [Bayesian Nash equilibrium] assumption that players *best-respond* to beliefs (*hence, we refer to these theories as belief-based*), but relax the requirement of consistency of beliefs". (p. 1435)

(emphasis ours).

[4]A broader, early approach is Stahl and Wilson (1994, 1995) who hypothesize a larger space of interesting types, including "worldly" types who respond to the true empirical distribution of play.

[5]This type was anticipated in research on signaling games, as discussed by Banks, Camerer and Porter (1994) and Brandts and Holt (1993). See also Binmore (1988).

Note that what ILN refer to as "belief-based models" combine two conceptually different features: nonequilibrium beliefs, and best response. But these modelling features can be combined in lots of ways. Indeed, the earliest forms of level-k and CH models (especially for empirical estimation) typically did assume some kind of quantal response, or interval of responses around a single best response (Nagel (1995); Stahl and Wilson (1994, 1995); Ho, Camerer and Weigelt (1998); Costa-Gomes, Crawford and Broseta (2001)).

Another class of models, quantal response equilibrium (QRE), retains the concept of accurate equilibrium beliefs (hence E), but relaxes the best response assumption (hence QR). Unlike cursed equilibrium, QRE was not specifically proposed to explain the winner's curse and associated phenomena which seem to be evidence of strategic naivete. However, it is often a useful benchmark model and can explain a wide variety of empirical deviations from best-response Bayesian-Nash equilibrium (e.g. Goeree, Holt and Palfrey (2002, 2005); Camerer (2003)).

Table 1 shows how different assumptions about accuracy of beliefs, and best response, can be combined more generally in different theories of strategic choice. In terms of this table, ILN use the phrase "belief-based models" to mean "best response models with nonequilibrium beliefs" (the lower left cell).

Table 1 also gives a count of the number of empirical papers analyzing nonequilibrium beliefs models which assume best-response (there are 24) and the number which assume quantal response (there are 37).[6] The counts indicate that a majority of studies of nonequilibrium beliefs assume quantal response. Thus, the models ILN call "belief based models", and reject in discussion of their data, are not those most commonly used in empirical practice.[7] This fact about model popularity, and the existence of the different combinations illustrated in Table 1, suggests two interesting empirical questions left unexplored by ILN: first, can an equilibrium quantal response model (QRE) fit the data? ILN conjecture (2010a) that "our evidence... is probably also evidence against QRE." We explore this conjecture further. Second, can quantal response forms of nonequilibrium belief models explain the data, and especially the empirical absence of treatment effects? We address this question by estimating

[6]These papers are listed in the references. The ones that assume best-response are [2-5], [8-9], [12], [17-18], [23], [27], [34-36], [42], [45-46], [57], [60-61], [64], [69], and [72-73]. The ones that assume quantal response are [1], [10], [13], [15-16], [19-21], [24], [28-33], [37-39], [41], [47-51], [62], [65-66], [68], [70], [74-78], and [80-82].

[7]In a brief discussion of the historical role of error in game theory, Goeree, Holt and Palfrey (2005) also note that forms of quantal response have played a prominent role in theory for many decades, in the form of trembles, and persistent and proper refinements of equilibria.

quantal response versions of both CE and CH models.

Table 1:

	Best Response	Quantal Response
Correct Beliefs	BNE	QRE
Nonequilibrium Beliefs	CE, Lk, CH (24)	QCE, QLk, QCH (37)

1.2 The Maximum Value Game

In the maximum-value auction of ILN, two players observe private signals x_1 and x_2 drawn independently from a uniform distribution over the set of integers, $\{0, 1, ..., 10\}$. They bid for an object which has a common value equal to the maximum of those two signals, $max(x_1, x_2)$. The highest bidder wins the auction and pays a price equal to the second-highest bid.[8].

How should you bid? Intuitively, if your own value is high you might underbid to save some money. And if your own value is low you might overbid because the object value, determined by the maximum of the two values, is likely to be higher than your low value.

These intuitions are not consistent with equilibrium, however. Bidding less than your signal is a mistake because your bid just determines whether you win, not what you pay, so bidding less never saves money. In fact, you might miss a chance to win at a price below its value if you bid too low and get outbid, so underbidding is weakly dominated. Second, bidding above your signal could be a mistake because if the other bidder is also overbidding, either of you may get stuck overpaying for an item with a low maximum-value. In the unique symmetric Bayesian-Nash equilibrium, therefore, players simply bid their values. In fact, the equilibrium can be solved for with two rounds of elimination of weakly dominated strategies.

While these bidding intuitions are mistakes, they can be small mistakes in terms of expected foregone payoffs. Moreover, payoff functions have many completely flat regions, so it turns out that some 'mistakes' of overbidding actually have zero expected cost in equilibrium. Whether mistakes are small or large is of central importance in the QRE approach. A key insight of quantal response equilibrium analysis is that even if large mistakes are rare, small

[8]Carrillo and Palfrey (forthcoming) study the maximum game with a bargaining mechanism.

mistakes could be common, and their equilibrium effects must be carefully analyzed. These equilibrium effects could be large, so small mistakes in payoff space can lead to large changes in the distribution of strategies. Therefore, even if over- and underbidding in maximum-value games are mistakes, those mistakes in bidding patterns could be common and could also have a substantial influence on bidding in general. Our paper explores the extent to which this possibility can explain the empirical results in ILN.

1.3 How Predictions Change with Quantal Response Combined with Belief Based Models

This section is designed to convey three ideas. First, we explain some intuitions of how bidding in phase I depends on the change from an assumption of *perfect* best response behavior (as in ILN) to *imperfect* best response, or quantal response, behavior. Second, we explain the intuition for the effect of imperfect beliefs, à la CE and level K models, when combined with imperfect best response behavior. Third, we provide a theoretical intuition for why the treatment effects (phase II and $MinBid$) can be negligible under both the imperfect best response models and the imperfect belief models.

First, to convey a sense of how QRE can approximate the behavior in phase I, it is useful to consider optimal and nearly optimal responses to the bidding behavior of your opponent, under different hypotheses about his or her degree of imperfect best response behavior. Recall that in the logit version of quantal response behavior, for each possible signal x, a bidder uses a behavioral strategy where the log probability of choosing each available bid is proportional to its expected payoff, where the proportionality factor, λ, can be interpreted as a responsiveness (or rationality) parameter. In particular, we have:

$$Pr(b|x) = \frac{exp(\lambda U(b|x))}{\sum_{a \in A} exp(\lambda U(b|x))}$$

where $b \in B$, the set of available bids, and $U(b|x)$ denotes the expected utility of bidding b conditional on observing x. The key question then is how a player forms his expectations, in particular, how a bidder's beliefs about the other bidder affects the expected utility term, $U(b|x)$.

At one extreme, if the hypothesis is that your opponent is bidding perfectly, in the sense of bidding according to the unique symmetric Nash equilibrium, then you believe he will simply bid his signal. If this is the case, then underbidding by you can lead to expected losses, however overbidding any amount is also a perfect best reponse, because if you win you will never pay more than your opponent's signal, which is less than or equal to the common value! Thus, in this case, which corresponds to behavior in a quantal response equilibrium with very high values of λ, we have $U(b|x) \approx U(b'|x)$ for all $x \in \{0, 1, ..., 10\}$ and for all $b, b' \geq x$. Thus, in a quantal response equilibrium, for high levels of rationality, payoff functions are completely flat on the entire bidding space bounded below by $b = x$.

Next, consider the opposite extreme belief about your opponent. Suppose the hypothesis is that your opponent is bidding completely randomly, corresponding to a bidder with a responsiveness parameter of $\lambda = 0$. If this is the case, then there is no correlation between your opponent's signal and his bid, so there is in fact no possibility of a winner's curse! The expected value of the object to you, conditional on winning, is simply $E[\max\{x_{you}, x_{opponent}\}|x_{you}] = \frac{x_{you}(x_{you}+1)}{11} + \frac{(10-x_{you})^2(11-x_{you})}{22}$. Straightforward calculations show that your optimal bid is equal to 5 for signals less than or equal to 5 and increases monotonically to 10 when your value is 10. Hence, the direct effect of quantal response behavior involves a combination of over- and underbidding, while the equilibrium effect of imperfect best responses leads to overbidding.

Figures 1a-c show three histograms of the predicted frequency of each bid under QRE, for three values of λ. The three values are the 25th ($\lambda = 0.114$), the 50th ($\lambda = 1.2975$), and the 75th ($\lambda = 4.509$) percentile of the distribution of λ_i estimated empirically for phase I of the Baseline and ShowBid treatments (as described later in Section 3). As λ increases from Figure 1a to Figure 1c, the distribution begins rather flat and then morphs to show more concentrated bids and substantial overbidding for all signals. The key transition is from Figure 1a to Figure 1b. When there is a lot of random bidding (from players with low λ values), it pays for low-signal bidders with higher λs to overbid, since they have a chance to get a bargain on high-value items if the high-signal players underbid. Similarly, high-signal players with higher λ's bid their signals or higher because underbidding is a substantial mistake in expected payoff terms. The highest value of λ gives a piecewise linear distribution that resembles an hockey-stick.

Of course, the discussion above simplifies the intuition of quantal response equilibrium; the

Figure 1: QRE Strategies for different λ's, Baseline & ShowBid Treatment, phase I

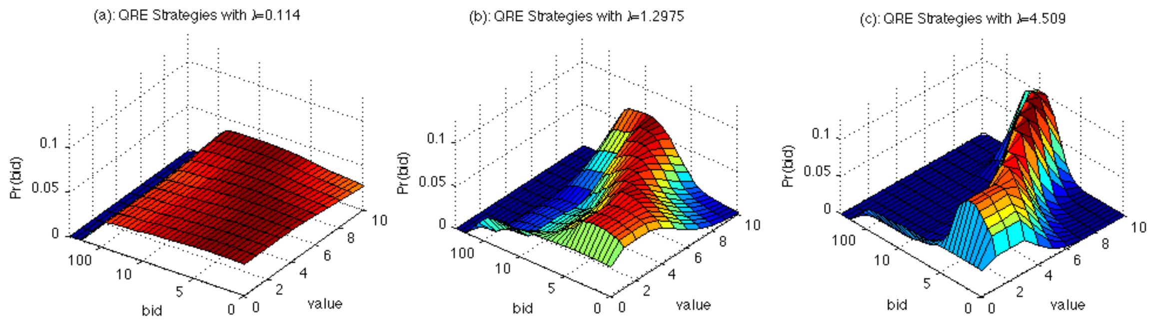

exact distribution of bids in a quantal response equilibrium depends in a complicated way on how all these equilibrium effects balance out. In fact, these interaction effects are even more complicated in the estimation of the next section, where we allow different bidders to have different response parameters, as in the heterogeneous version of QRE (HQRE) developed in Rogers, Camerer and Palfrey (2009).

The effects of CE and CH complement, and in some cases reinforce, the effects of imperfect best response. With respect to CE, the easiest way to see this is to note that if $\lambda = 0$ there is no winner's curse, so imperfect best responses implies some degree of partially cursed behavior. Both effects (CE and QRE) push bids up for low signal values, and reduce to expected losses from overbidding. In level-k and CH models, Level 1 bidders behave exactly as if they are facing $\lambda = 0$, so the discussion of optimal responses is the same as described above. However, level-2 bidders who think they are playing level-1 (the level-k specification, not CH) know they are playing overbidders, and therefore they might mistakenly pay too much, so their optimal response is to bid their signal. To consider higher levels, recall that if an opponent is bidding their signal, then overbidding and bidding one's signal are both weak best-responses. Therefore, a level-3 bidder is actually expected to overbid, in a statistical sense, against a level-2 bidder. Following a similar logic, higher level players alternate from level to level between overbidding (k even) and signal-bidding (k odd). Thus, the level-k model assuming best response predicts a mixture of random bidding, overbidding, and signal-bidding, as in QRE.

What are the implications of these models of imperfect best response and imperfect beliefs for Phase II bidding behavior? In Phase II, each bidder competes against a machine player that is programmed to bid exactly according to the bidding function the bidder used in Phase

1. Because there is no feedback between phases, in our analysis we assume that the degree of responsiveness or cursedness is the same in Phase I and Phase II, but take into account that their opponent's bidding strategy is governed by a different distribution in Phase I and Phase II.[9] First, consder the QRE model. In this case, if a bidder is a low-λ_i bidder who is relatively unresponsive to payoff differences, we would expect a lot of variance in the form of under- and over-bidding in Phase II as well as Phase I. Since the bidders who grossly overbid in the Phase I reveal themselves to have very low λ this kind of behavior is expected to persist in Phase II – indeed that is what the data shows. By the same token, bidders who were bidding close to their signal in Phase I should bid close to their signal in Phase II. For intermediate values of λ one also expects relatively small changes in behavior, if any. Thus QRE predicts rather small treatment effects.

What about the CE model in phase II? While the players know the bidding strategy of the computerized opponent, they must bid without knowing the signal drawn by the opponent in a specific round. Thus, one can still apply the cursed equilibrium model, where one of the players (the computer) simply has a fixed strategy. And, as in CE, there is no reason to expect a human bidder who fell prey to cursed beliefs in Phase I would necessarily have an epiphany and suddenly realize how to correctly account for the correlation between the machine bidder's signals and bids. As noted in our introduction, behavior based on cursed beliefs has been widely observed in decision experiments with human players competing against computerized rules in situations that are isomorphic to games. Thus, in Phase II, as in Phase I, cursed behavior combined with QRE could produce a relatively modest increase in the extent of overbidding relative to a pure QRE model.

Finally, consider the implications of level K models for Phase II. As ILN point out, with a *perfect* best response model, a level k bidder in Phase I should behave like a level k+1 bidder in Phase II. For example, if a bidder is a value bidder (k odd) in Phase I he should be an overbidder in Phase II and vice versa. This is not what one observes in the data. However, the picture changes in two interesting ways if there is quantal response. First, bidders who are level-0 or odd-level bidders in phase I (because they overbid) could be revealing a low response sensitivity λ_i. For example, their (stochastic) bidding function could look like Figure 1a or 1b above. If λ_i is not too large, they might respond very little to knowing that they are bidding against their own previous phase I bids. Second, bidders who bid close to their

[9] Subjects were given zero feedback between Phase I and Phase II so there is no opportunity to 'unlearn' cursed beliefs or reduce or increase their payoff responsiveness.

signals in phase I will be identified as candidate even-level bidders. But when they are bidding against their own previous bids in phase I (which are bids equal to signals), then bidding signals and bidding above signals are *equally "best"* responses. Even a little quantal response implies that these bidders may actually bid above their signals, and hence above their own previous bids.[10]

So the combined effect of quantal response on the analysis of the response to the phase II treatment is to dampen (depending on λ) the predicted effect of the level 1's bidding lower, and to add an effect of level 2's bidding *higher*. Thus, it is not clear there will be any strong effect overall in aggregate bids in the comparison of phase I to phase II. The same conclusion holds for the CH model with quantal response, which predicts identical level-1 behavior and a similar level-2 response to the phase II treatment.

2 Experimental Design Details

In the experiment, signals are said to be uniformly, independently, and discretely distributed from 0 to 10. Subjects make a series of bids but do not receive any feedback after each auction. In fact, each subject received each integer value exactly once; so the resulting bid function is similar to what you might get if you asked for a bid for each possible value (a.k.a. the strategy method). Payoffs are $.50 per experimental dollar. Subjects begin with $20 and received that sum plus or minus the total of all the payoffs from random matching with other bidders. Notice that it is possible to lose money, if you overbid persistently. In fact, about 27.50% of subjects lost money from their original base pay of $20 and 16.88% ended up with a negative net balance (which they did not have to pay).

In a pilot experiment, ILN allowed subjects to bid above the maximum value, up to $100, and they noticed substantial overbidding. They therefore decided to allow subjects in the main experiment to bid up to $1 million.[11]

[10] Indeed, this shift does seem to happen: ILN classify 20 players as signal-bidders in phase I. Of those, 13 (65%) are also classified as signal-bidders in phase II, but 6 (20%) are classified as overbidders (some even bidding above the maximum value of 10).

[11] The instructions and text are not crystal clear on how the maximum bids are implemented. The instructions appended to the paper do not say anything about maximum bid caps, but the paper is clear about the maximum.

3 Experimental Results and New Analysis

3.1 Preliminary Analysis of Experimental Data

Let's start with a look at all the overall bids.[12] Figure 2 reports the bids by signal for phase I of the *Baseline* and *ShowBid* treatments (the two treatments are pooled because the results are very similar). In the left graph, all bids are shown on a log scale (bids are between \$0 and \$1,000,000). To gauge the extent of overbidding, we add to this scatterplot two lines: a continuous line drawn at a level of bidding equal to the private signal received, and a dashed line signaling a bid equal to the upper bound of the common value (i.e. \$10). The left graph reports all bids. The right graph, using a linear scale, truncates bids above \$15 and reports only the bids between \$0 and \$15. The continuous line in this graph is the median bid as a function of the private signal (considering all bids).

An eye-catching feature of the Figure 2 scatterplot of all the bids is the bids which are very high (relative to the maximum object value of \$10). Overall, 7.66% of bids are \$100 or above, and 2.27% of the bids are exactly \$1,000,000. Note that these are actual (half) US dollars (converted at \$.50 per experimental dollar bid to payment)– an experimental bid of \$1,000,000 therefore corresponds to a subject imagining paying \$500,000 in actual dollars. These high bids are present in all sessions and treatments. We discuss them briefly in the conclusion.

Excluding bids above \$15, there is substantial evidence of both underbidding at high values, and overbidding throughout (but especially for lower values). The median bids are relatively flat, at around \$5, for values \$0-\$5, then increase monotonically for higher values (the "hockey stick" distribution).

Figure 2 reports the same data for phase II of the *Baseline* and *ShowBid* treatments, while Figure 3 is for phase I of the *MinBid* treatment. Except for the truncation of bids below values in MinBid, the pattern of bids in phase I and phase II look very similar.

[12] ILN report their results quite compactly, in the form of a heuristic classification of bidding strategies.

Figure 2: Bids by Signal, Baseline & ShowBid Treatment, phase I

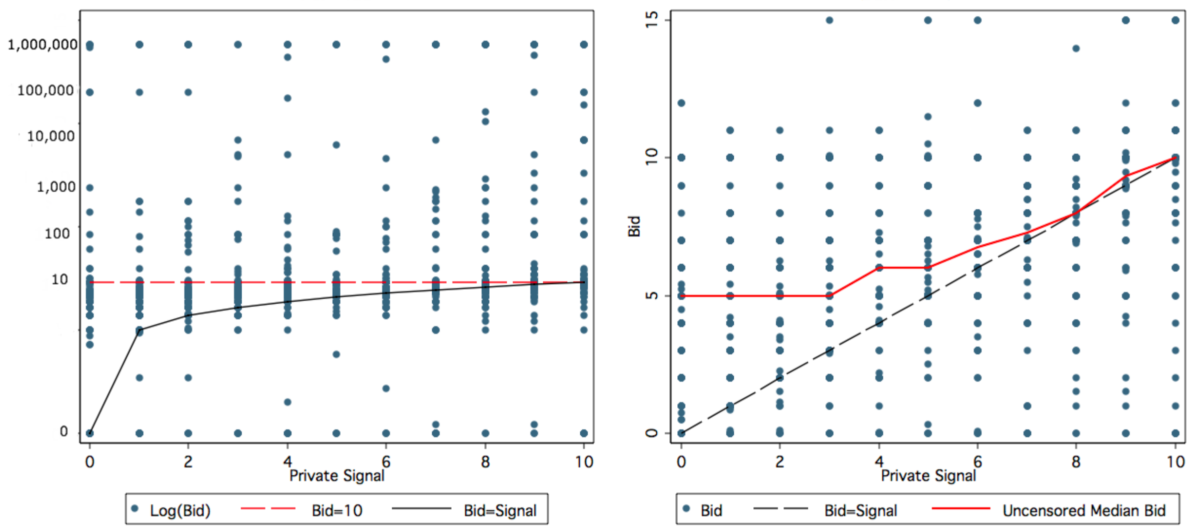

Figure 3: Bids by Signal, Baseline & ShowBid Treatment, phase II

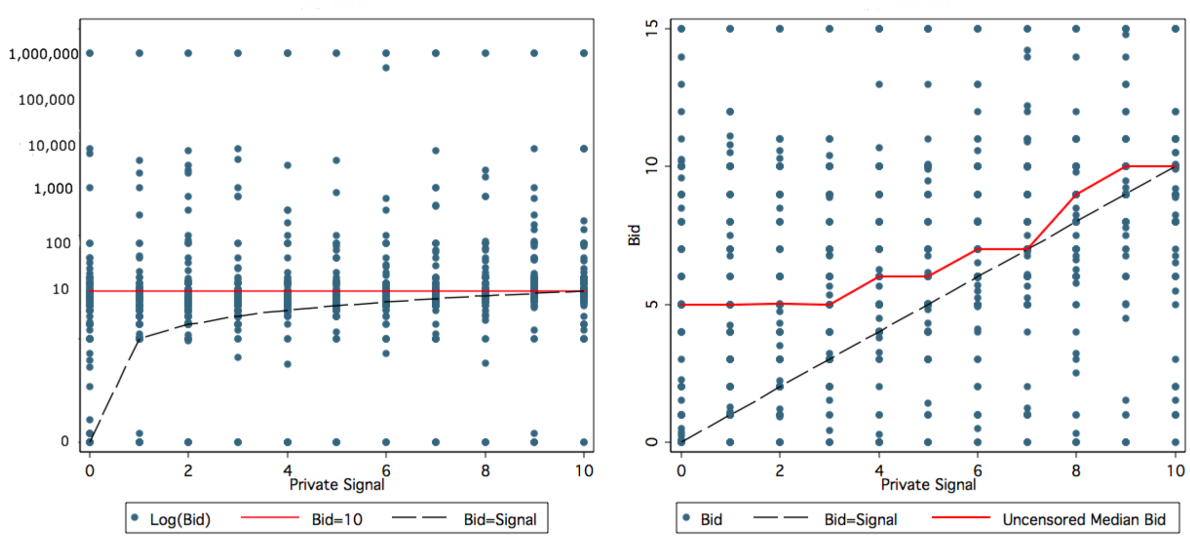

Figure 4: Bids by Signal, MinBid Treatment, phase I

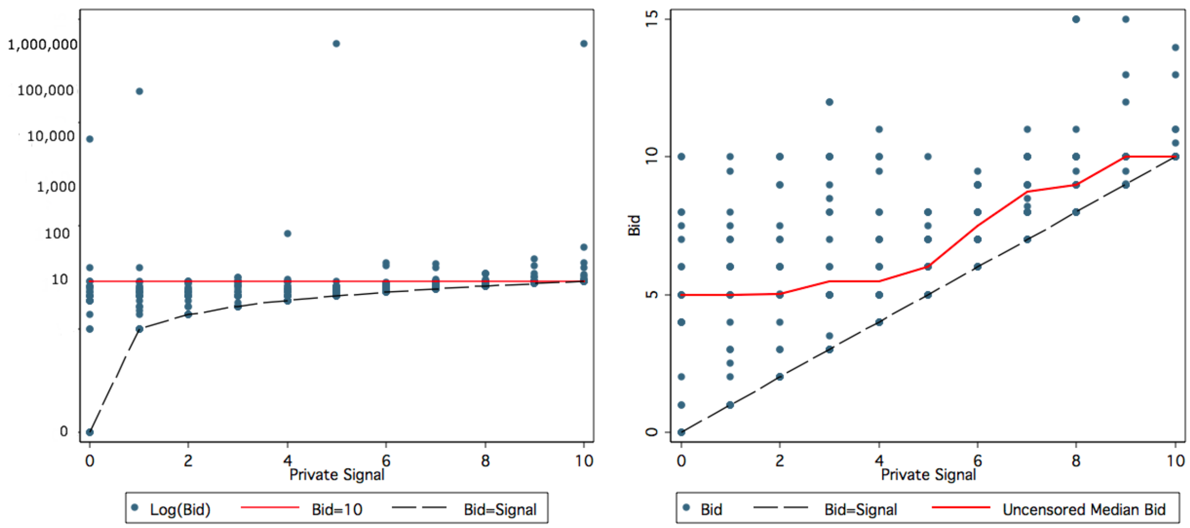

3.2 Model Estimation and Forecasting

3.2.1 HQRE Predictions

In the logit version of HQRE (Rogers, Camerer and Palfrey (2009)) each agent, at each information set, uses a behavioral strategy where the log probability of choosing each available action is proportional to its expected payoff, where the proportionality factor, λ_i, can be interpreted as a responsiveness (or rationality) parameter that can vary between subjects. In particular, we have:

$$Pr(a) = \frac{exp(\lambda_i U(a))}{\sum_{a \in A} exp(\lambda_i U(a))} \tag{1}$$

where $a \in A$ is one of the actions available to the agent.

In the maximal value auction, when agent i receives private signal s_i, his expected utility from bid b_i is the following:

$$EU(b_i|s_i) = Pr(b_i > b_j)\left[E(max(s_i, s_j)|b_i > b_j) - E(b_j|b_j < b_i)\right] + \frac{1}{2}Pr(b_i = b_j)\left[E(max(s_i, s_j)|b_i = b_j) - b_i\right]$$

The first term is the probability of winning times the expected net benefit conditional on winning, while the second term is the probability of a tie times the expected net benefit conditional on tying (ties are broken randomly), and given the strategy of the other bidder. Given these expected payoffs, the choice probabilities follow from the choice rule in (1), i.e.:

$$Pr(b_i|s_i) = \frac{exp(\lambda_i EU(b_i|s_i))}{\sum_{b \in B} exp(\lambda_i EU(b|s_i))} \tag{2}$$

An insight from Bajari and Hortacsu (2005) enables us to estimate HQRE consistently. They note that if the model is correct, then the actual data provide an unbiased estimator of the aggregate joint distribution of bids and signals – and also players' beliefs, since beliefs are assumed to be consistent with actual behavior, by the E part of the definition of QRE. Therefore, we use the empirical joint distributions of signals and bids (i.e. the distributions observed in the experiments) in phase I of the *Baseline* and *ShowBid* treatments to compute the empirical expected utility of each bid given each private signal, $\widehat{EU}(b_i|s_i)$.[13] Then, for each subject playing in those sessions, we estimate, using MLE, the $\widehat{\lambda}_i$ that best fits his observed bidding function. This gives us a distribution over bids (conditional on the signal) for each subject. Aggregating over all subjects, we obtain a predicted distribution of bids for each signal.

In Panel 1 of Figure 4 we show the quartiles of this distribution and we compare them with the quartiles of the observed bids (represented with boxplots). The predictions for phase II (Panel 2 of Figure 4) are, instead, out-of-sample predictions. For each subject, we first compute $\widehat{EU}(b_i|s_i)$ using the empirical joint distribution of bids and signals from his previous plays in phase I (i.e. each agent plays against his exact previous bidding function). Then, we predict his behavior in phase II by plugging in these expected payoffs and the $\widehat{\lambda}_i$ estimated

[13]For tractability, we bin bids for the estimation. The bins use \$.50 intervals for bids less than or equal to \$12, with progressively coarser bins for higher values.

for phase I into (2).

3.2.2 CE-QRE Predictions

An agent i who is *cursed*, in the sense of Eyster and Rabin (2005), correctly predicts the distribution of other players' bids, but underestimates the degree to which these bids are correlated with the other players' private signals. In particular, the expected utility of a *cursed* agent is given by:

$$\begin{aligned} EU_{CE}(b_i|s_i) &= Pr(b_i > b_j)[E(max(s_i, s_j) - E(b_j|b_j < b_i)] \\ -\frac{1}{2}Pr(b_i &= b_j)[E(max(s_i, s_j) - b_i] \end{aligned}$$

The extent to which agents are *cursed* is parameterized by the probability $\chi \in [0, 1]$ they assign to other players playing their average distribution of actions irrespective of type rather than their type-contingent strategy (to which she assigns probability $1 - \chi$). In the maximal value game, the expected utility of a χ-cursed agent i from bidding b_i when observing signal s_i is given by:

$$EU_{CE}^{\chi}(b_i|s_i) = (1 - \chi)EU(b_i|s_i) + \chi EU_{CE}(b_i|s_i)$$

We assume, moreover, that agents respond imperfectly as in a HQRE:

$$Pr(b_i|s_i) = \frac{exp(\lambda_i EU_{CE}^{\chi}(b_i|s_i))}{\sum_{b \in B} exp(\lambda_i EU_{CE}^{\chi}(b|s_i))}$$

As in the HQRE estimation, we use the empirical joint distributions of signals and bids in phase I of the *Baseline* and *ShowBid* treatments to compute $\widehat{EU}_{CE}^{\chi}(b_i|s_i)$. Then, we estimate with a standard MLE procedure the vector of individual $\widehat{\lambda_i}$'s and the cursed parameter $\widehat{\chi}$ that fit the data best. To compute the out-of-sample predictions for phase II we use the same $\widehat{\lambda_i}$'s and $\widehat{\chi}$ estimated for phase I but we take into account that everybody has different beliefs (based on the empirical distribution of their previous bids in phase I) and, thus, different expected payoffs. The predicted distributions for phase I are shown in Panel 1 of Figure 5, while the out-of-sample predictions for phase II are in Panel 2 (assuming the same χ value).

Intuitively, CE generates overbidding in phase I because players ignore, to an extent calibrated by χ, the connection between the bid and signal of the other player. Therefore, bidders with low signals do not realize that a low bid indicates the other player's signal may be low, and hence overbid to take advantage of the perceived bargain.

3.2.3 CH-QRE Predictions

In the CH model (Camerer, Ho and Chong (2004)) players adopt an iterative decision rule based on k steps of thinking, and assume a Poisson frequency distribution $f(k)$ of step k players. "Step 0" types do not assume anything about their opponents and merely choose according to a uniform probability distribution. "Step k" players assume that their opponents are distributed, according to a normalized Poisson distribution, from step 0 to step "$k-1$". Therefore, they accurately predict the relative frequencies of less sophisticated players, but ignore the possibility that some players may be as sophisticated or more. Step 1 players assume that all the others players are step 0 players, while step 2 players assume the other players are a combination of step 0 players and (quantal responding) step 1 players. In particular, a step k player's belief (for $k > 0$) about the proportion of h step players is given by $f_k(h) = 0$ for $h > k-1$ and $f_k(h)$ for $h < k$, given by:

$$f_k(h) = \frac{f(h)}{\sum_{n=0}^{h} f(k)}$$

where $f(k) = e^{-\tau}\tau^k/k!$, which is characterized by one parameter $\tau > 0$.

In this setting, denote the expected utility of step 1 players as $EU_{CH_1}(b_i|s_i,\tau)$, and the one of step 2 players as $EU_{CH_2}(b_i|s_i,\tau)$. We assume, moreover, that agents respond imperfectly as in a QRE with a homogeneous λ:[14]

[14] In principle, one could allow for heterogeneous λ in the CH model. However, the CH model already has built into it different types so we conjecture that allowing for heterogeneity in responsiveness would add little improvement in fit. Thus, for parsimony we estimate only a homogeneous CH-QRE model. Also, to keep the estimation simple, we only consider types $k < 3$. Allowing for higher levels in the estimation adds little.

$$Pr_k(b_i|s_i) = \frac{exp(\lambda EU_{CH_k}(b_i|s_i))}{\sum_{b \in B} exp(\lambda EU_{CH_k}(b|s_i))} \tag{3}$$

For each pair (λ, τ), (3) implies a joint distribution of bids and signals for type 1 and type 2 players (and, combining this with the random behavior of type 0 players, an aggregate distribution of bids). To generate the predictions in Panel 1 of Figure 6, we estimate with MLE the pair ($\widehat{\lambda}$, $\widehat{\tau}$) that fits best the empirical strategies from phase I of the *Baseline* and *ShowBid* treatments. For the out-of-sample predictions for phase II we use the same $\widehat{\lambda}$ estimated for phase I and compute the corresponding QRE bidding strategies. Notice that the estimated parameter for the Poisson distribution of types, $\widehat{\tau}$, has no role in phase II, as players are now bidding against their own previous bid profile rather than against a mixture of lower level players. Moreover, even if $\widehat{\lambda}$ is the same for every agent, we take into account that each bidder has different beliefs (based on the empirical distribution of her previous bids in phase I) and, thus, different expected payoffs. We present out-of-sample predictions for phase II of the same treatments in Panel 2 of Figure 6.

We present two sets of CH-QRE predictions for the *MinBid* treatment. The left-hand panel of Figure 7 shows predictions for phase I obtained by estimating ($\widehat{\lambda}$, $\widehat{\tau}$) as described above. The right-hand panel of Figure 7, instead, shows predictions for phase I of the *MinBid* treatment obtained using the parameters estimated with the data from phase I of the *Baseline* and *ShowBid* treatments.

4 General Discussion and Conclusion

The interesting empirical results originally reported by ILN in the maximum-value auction show a hockey stick-shaped pattern of overbidding: bids for signals from 0-5 are around 4-6; higher signals also exhibit modest overbidding that increases linearly with the signal. Models with heterogeneous quantal response fit these data rather well. Those models fit the central hockey-stick tendency; and they also generate approximately the right degree of variation in bids, measured by interquartile range. The fit is further improved (significantly) by weakening the assumption of equilibrium beliefs (in QRE) according to the one-added-

Figure 5: HQRE Predictions, Baseline & ShowBid Treatment

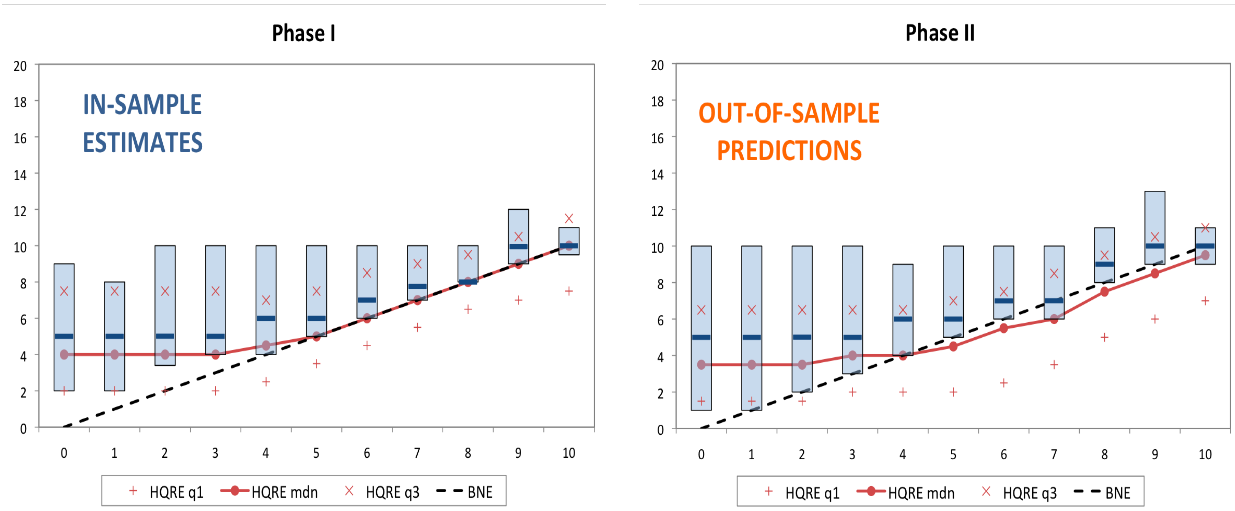

Figure 6: CE-QRE Predictions, Baseline & ShowBid Treatment

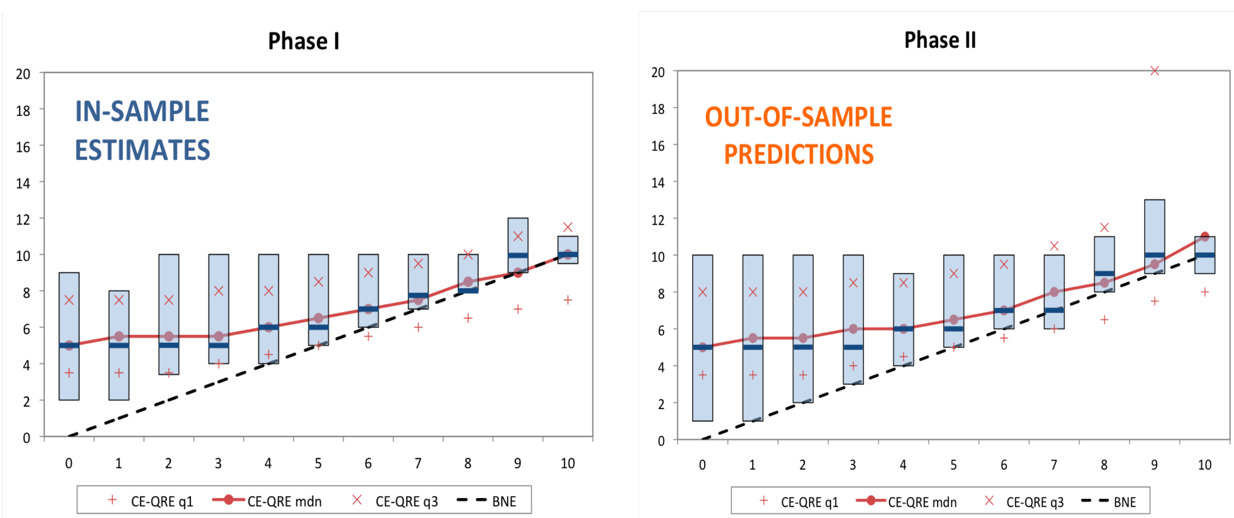

Figure 7: CH-QRE Predictions, Baseline & ShowBid Treatment

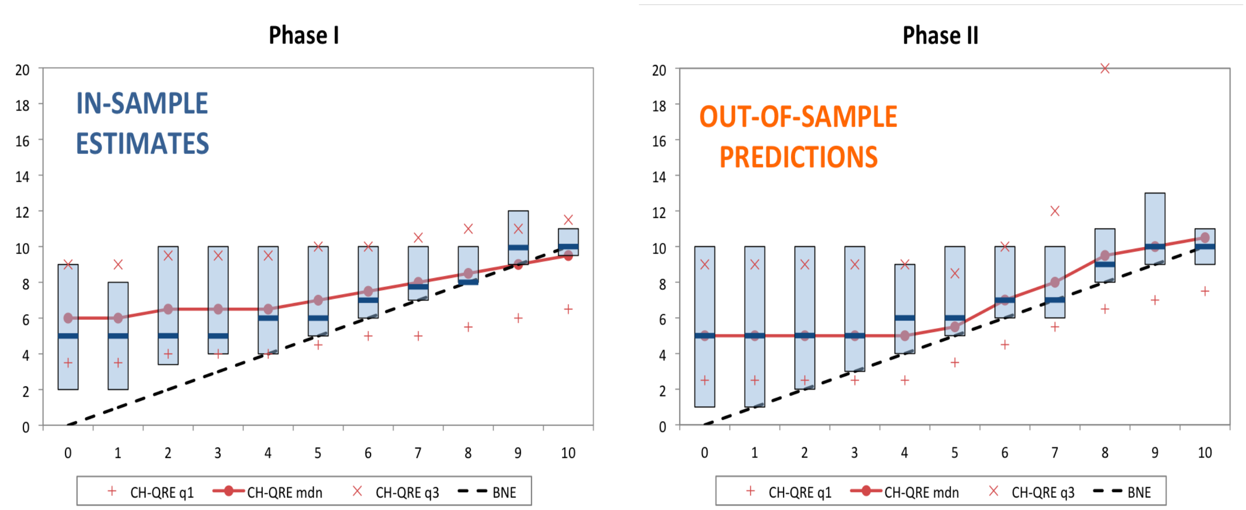

Figure 8: CH-QRE Predictions, MinBid Treatment, phase I

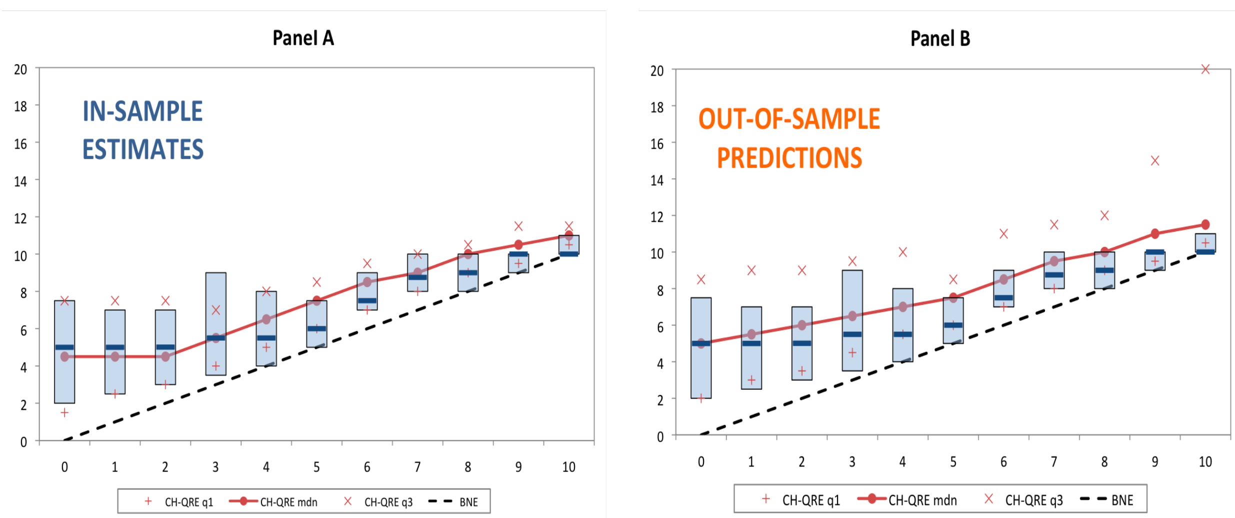

Table 2: Estimated Parameters

	Baseline	**MinBid**
# individuals	108	26
# observations	1188	286
HQRE		
$\widehat{\lambda}_i$ - 1st quartile	0.114	0.883
$\widehat{\lambda}_i$ - 2nd quartile	1.2975	1.748
$\widehat{\lambda}_i$ - 3rd quartile	4.509	4.083
-Log-likelihood	3590	912
CE-HQRE		
$\widehat{\chi}$	1	1
$\widehat{\lambda}_i$ - 1st quartile	0.1065	1.565
$\widehat{\lambda}_i$ - 2nd quartile	1.442	2.9925
$\widehat{\lambda}_i$ - 3rd quartile	4.3105	4.272
-Log-likelihood	3580	708
CH-QRE		
$\widehat{\tau}$	0.9	3.4
$\widehat{\lambda}$	2.01	13
-Log-likelihood	3730	759

parameter versions of CE or CH. The addition of those nonequilibrium CE or CH beliefs raises the level of predicted bids, which are about $1 too low in the QRE fits, to match the general level of bidding more closely.

More importantly, the estimates obtained from the phase I data fit the phase II data out-of-sample about as well as they fit phase I data in-sample, and fit the MinBid treatment similarly well. This means that the treatment effects which were hypothesized by ILN based on perfect best-response models, are *not* hypothesized to occur under quantal response. Those hypothesized changes are not evident in the data either.

In their original analysis of these data, ILN draw a different broad conclusion than we do. The difference is substantive, not semantic, and requires a careful distinction between two different dimensions by which models may deviate from standard Nash equilibrium: perfect vs. imperfect best response; and equilibrium ("rational expectations") vs. nonequilibrum beliefs.

ILN's expressed interest in this bidding game is that some specialized *perfect best-response* models of nonequilibrium beliefs (CE, and level-k) predict that bids will be lower under the two experimental treatments, for certain types of bidders. The idea is that people may overbid in phase I because they do not correctly infer the connection between signals and bids (CE), or because they think they are facing random bidders (level-1). In theory, these erroneous beliefs should be erased and replaced in phase II when they bid against their own previous bid functions. The phase II treatment is therefore predicted to make their bids lower. Since, empirically, bids are *not* much lower in the treatment conditions, ILN conclude, "Overall, our study casts a serious doubt on theories that posit the WC is driven by beliefs."(p. 1435)

However, that conclusion must be qualified by "...assuming perfect best responses" at the end (or implicitly qualified by a reader correctly noticing that "driven by beliefs" is used as shorthand to mean nonequilibrium belief modeling with perfect best response in the definition).[15] Virtually the opposite conclusion is reached if one combines nonequilibrium belief models with quantal response. Technically, both conclusions are correct; and together they provide a full picture of what the data show about different kinds of theories. ILN's analysis shows that best response models clearly fail in this environment.[16] We show that models

[15]This important qualification was later expressed by ILN coauthor Niederle (2010) who describes the ILN results and says "The results of the experiment are in general bad news for any theory that relies on best response behavior" (p.12).

[16]The same conclusion was reached by Carrillo and Palfrey (forthcoming) in their independent contempo-

based on quantal response equilibrium do fit the data well, and even better when combined with nonequilibrium beliefs.

The rest of our conclusion consists of four remarks.

First, note that nothing about these data or these models necessarily guarantees a good fit for HQRE. In fact, the approximate accuracy of HQRE in fitting the overall distribution, in both baseline and treatment conditions, came as a small surprise to us. It might surprise others too. For example, while the published ILN paper (2010b) does not mention QRE at all, the penultimate draft (2010a) says: "Informally, our evidence against (a') [noisy best response with zero median in noise] in initial periods of play in common-value auctions is probably also evidence against QRE. However, formally, given that we cannot test assumption (a') [noisy best response], we cannot say much about QRE."

Their first sentence is not right because QRE generally does *not* predict zero median in noise – as this specific game beautifully illustrates. Noise around equilbrium strategies is driven by expected payoff differentials, which in this game are highly *asymmetric* around signal bidding; so probable evidence against a model with zero median noise could be evidence *for* QRE, and is certainly not clear evidence against it. Also, our paper does show one way to *can* "test assumption a'", by using the maintained assumption of equilibrium beliefs to justify using the data as an estimator of beliefs. An earlier working paper (ILN 2008) says more about QRE, but we do not comment further on the grounds that authors should not be held responsible for all details of early paper drafts.

Second, the maximum-value auction and associated treatments are not that useful for comparing different theories with equilibrium or nonequilibrium beliefs, under best response. All those theories are clearly rejected by the phase I behavior alone. In contrast however, the game is actually quite useful for comparing models with best response and quantal response. Best response models cannot explain the distribution of bids and lack of treatment effects, while quantal response models can. Given that best response and quantal response models are both used frequently along with nonequiibrium beliefs (as shown by the Table 1 counts), it is handy to have a game that can sharply distinguish them.

Third, the prevalence of very high bids in these auctions is striking: 2% of $1 million; and 8% of $100 or more. In their design, a $1 million experimental bid converts to $500,000 US dollars in subject payments. These sky-high bids are evident in all sessions and treatments.

raneous study of the maximum value game in a bargaining setting.

Many of these bids are concentrated in a few individuals, but 27% of subjects (36 out of 134) make at least one bid of $100 or above, which is ten times the maximum object value. ILN do not discuss or explain these interesting high bids at all. The published paper implicitly acknowledges existence of bids above the theoretical maximum of 10, stating (footnote 32, p. 1443): "We plot median, rather than average, bids because averages are distorted by bids above 10."[17]

One hypothesis is that subjects were amused by the possibility of actually bidding $1 million, knowing full well they wouldn't have to pay if they lost even a small fraction of that amount. On the other hand, such bids are far from irrational. In fact, an alternative hypothesis is that they are playing in accordance with one of the infinite number of asymmetric "bullying" Nash equilibria in which a small percentage of subjects bid above 10, and other subjects weakly best respond by bidding zero (which increases the bully's payoff). ILN note that nothing in the experimental design encourages emergence of such asymmetric equilibria[18]; on the other hand, there is little to discourage it since there is no feedback to unlearn behavior if it is out of equilibrium.

Fourth, the maximum-value auction results provide an opportunity to reflect upon how much we demand, expect, and hope that models of human behavior can do. It is certainly the case that the winner's curse in common-value strategic auctions, and similar mistakes in single-person decision analogues (the "Beastie Run" example from 1980's decision research, and Charness and Levin, (2009)) and common value bargaining settings (Carrillo and Palfrey, forthcoming), clearly point to a widespread failure to correctly compute - or perhaps even comprehend - conditional expectations. However, it is also true that simple models of the difficulty of computing conditional expectations (as well as CE) cannot explain deviations from equilibrium in complete information games or in Bayesian games with a private values structure.[19] An advantage of level-k and CH and QRE explanations is that they can explain anomalies in *both* complete information games *and* in private information games like the maximum value auction (as results reported here demonstrate)– because limits on the

[17]Their 2008 version gives a little more information (footnote 39, p. 16): "We plot median, rather than average, bids because averages are distorted by bids above 10 (which are sometimes very high, even up to 1000000)."

[18]ILN state: "In our experiment, matching of subjects is anonymous and there is no feedback, so that asymmetric Nash Equilibria do not seem plausible." (footnote 11, p. 7, 2008)

[19]Analogy Based Equilibrium (Jehiel (2005) and Jehiel and Koessler (2008)) is somewhat more general and can have interesting implications in some games with private values.

depth of strategic thinking about others automatically generate mistaken inferences about information-strategy links (cf. Brocas et al. (2010)). Probably the best collection of modelling tools will include quantal response (as our results in this paper suggest), level-k or CH, and some specialized account of why conditioning expectations on hypothetical events that must be imagined (even in apparently simple decisions) is so tricky. Progress beyond the useful benchmark of Nash equilibrium has been so rapid in recent years, however, that it is easy to be optimistic about making further progress on a more integrated, predictive general theory that can be tractably applied to the empirical analysis of all these types of games and decisions.

References

[1] Agranov, Marina, Andew Caplin, and Chloe Tergiman. 2010. " The Process of Choice in Guessing Games." Unpublished

[2] An, Yonghong. 2010. "Nonparametric Identification and Estimation of Level-k Auctions". Unpublished

[3] Arad, Ayala. 2010. "The Tennis Coach Problem: A Game-Theoretic and Experimental Study." Unpublished

[4] Arad, Ayala, and Ariel Rubinstein. 2010. "The 11-20 Money Request Game: Evaluating the Upper Bound of k-Level Reasoning." Unpublished.

[5] Aradillas-Lopez, Andres, and Elie Tamer. 2008. "The Identification Power of Equilibrium in Simple Games." Journal of Business & Economic Statistics, 26(3): 261-283.

[6] Bajari, Patrick, and Ali Hortacsu. 2005. "Are Structural Estimates of Auction Models Reasonable? Evidence from Experimental Data." Journal of Political Economy, 113(4):703-741.

[7] Banks, Jeffrey, Colin Camerer and David Porter. 1994. An Experimental Analysis of Nash Refinements in Signaling Games. Games and Economic Behavior, 6(1):1-31.

[8] Bardsley, Nicholas, Judith Mehta, Chris Starmer, and Robert Sugden. 2010. "Explaining Focal Points: Cognitive Hierarchy Theory "versus" Team Reasoning." Economic Journal, 120(2), 40:79.

[9] Berger, Ulrich, Hannelore De Silva, and Gerlinde Fellner. 2010. "Cognitive Hierarchies in the Minimizer Game." Unpublished

[10] Bhatta, Meghana A., Terry Lohrenz, Colin F. Camerer, and P. Read Montague. 2010. "Neural signatures of strategic types in a two-person bargaining game." Proceedings of the National Academy of Sciences, 107(46), 19720:19725.

[11] Binmore, Ken. 1988. "Modeling Rational Players: Part II." Economics and Philosophy 4(1):9-55.

[12] Bosch-Domenech, Antoni, Jose Garcia-Montalvo, Rosemarie Nagel, and Albert Satorra. 2002. "One, Two, (Three), Infinity...: Newspaper and Lab Beauty-Contest Experiments." American Economic Review, 92(5), 1687-1701.

[13] Bosch-Domenech, Antoni, Jose' G. Montalvo, Rosemarie Nagel, and Albert Satorra. 2010. "A finite mixture analysis of beauty-contest data using generalized beta distributions." Experimental Economics, 13(4) 461:475.

[14] Brandts, Jordi and Charles A. Holt. 1993. "Adjustment Patterns and Equilibrium Selection in Experimental Signaling Games." International Journal of Game Theory, 22(3):279-302.

[15] Breitmoser, Yves. 2010. "Hierarchical Reasoning versus Iterated Reasoning in p-Beauty Contest Guessing Games." Unpublished.

[16] Brocas, Isabelle, Juan D. Carrillo, Colin F. Camerer, and Stephanie W. Wang. 2010. Imperfect choice or imperfect attention? Understanding strategic thinking in private information games." Unpublished.

[17] Brocas, Isabelle, Juan D. Carrillo, and Manuel Castro. 2010. "The Nature of Information and its Effects on Bidding Behavior: Laboratory Evidence in a Common Value Auction." Unpublished

[18] Bottero, Margherita. 2010. "Cognitive Hierarchies and the Centipede Game". Unpublished

[19] Brown, Alexander, Colin F. Camerer, and Dan Lovallo. 2010. "To Review or Not To Review? Limited Strategic Thinking at the Movie Box Office." Unpublished.

[20] Burchardi, Konrad B., and Stefan P. Penczynski. 2010. "Out Of Your Mind: Estimating The Level-k Model." Unpublished

[21] Cai, Hongbin, and Joseph Tao-Yi Wang. 2006. "Overcommunication in Strategic Information Transmission Games." Games and Economic Behavior, 56(1): 7:36.

[22] Camerer, Colin. 2003. *Behavioral Game Theory.* Russell Sage Foundation, New York, New York/Princeton University Press, Princeton, New Jersey.

[23] Camerer, Colin, Ho, Teck Hua, and Chong, Juin-Kuan. 2004. "A Cognitive Hierarchy Model of Games." Quarterly Journal of Economics, 119(3):861:898.

[24] Carrillo, Juan D., and Thomas R. Palfrey. 2009. "The Compromise Game: Two-Sided Adverse Selection in the Laboratory." American Economic Journal: Microeconomics, 1(1): 151:181.

[25] Carrillo, Juan D., and Thomas R. Palfrey. Forthcoming. "No Trade." Games and Economic Behavior.

[26] Carroll, John S., Max H. Bazerman, and Robin Maury. 1988. "Negotiator cognitions: a descriptive approach to negotiators' understanding of their opponents." Organizational Behavior and Human Decision Processes. 41, 352-370.

[27] Charness, Gary and Dan Levin. 2009. "The Origin of the Winner's Curse: A Laboratory Study." American Economic Journal: Microeconomics, 1(1): 207:236.

[28] Chong, Juin Kuan, Colin F. Camerer, and Teck-Hua Ho. 2005. "Cognitive Hierarchy: A Limited Thinking Theory in Games." In Rami Zwick and Amnon Rapoport (editors), Experimental Business Research: Marketing, Accounting and Cognitive Perspectives. Boston: Kluwer Academic Press.

[29] Costa-Gomes, Miguel A., and Vincent P. Crawford. 2006. "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study." American Economic Review, 96(5): 1737-1768.

[30] Costa-Gomes, Miguel A. and Vincent P. Crawford. 2007. "Studying Cognition via Information Search in Two-Person Guessing Game Experiments." Unpublished

[31] Costa-Gomes, Miguel A., Vincent P. Crawford, and Bruno Broseta. 2001. "Cognition and Behavior in Normal-Form Games: An Experimental Study." Econometrica, 69(5): 1193-1235.

[32] Costa-Gomes, Miguel A., Vincent P. Crawford, and Nagore Iriberri. 2009. "Comparing Models of Strategic Thinking in Van Huyck, Battalio, and Beil's Coordination Games." Journal of the European Economic Association, 7(2-3): 365:376.

[33] Costa-Gomes, Miguel A., and Georg Weizsecker. 2008. "Stated Beliefs and Play in Normal-Form Games." Review of Economic Studies, 75(3): 729-762.

[34] Crawford, Vincent P. 2003. "Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions." American Economic Review, 93(1): 133-149.

[35] Crawford, Vincent P. 2007. "Let's Talk It Over: Coordination via Preplay Communication with Level-k Thinking." Presented at the Arne Ryde Symposium on Communication in Games and Experiments. Unpublished.

[36] Crawford, Vincent P., Uri Gneezy, and Yuval Rottenstreich. 2008. "The Power of Focal Points is Limited: Even Minute Payoff Asymmetry May Yield Large Coordination Failures." American Economic Review, 98(4): 1443:1458.

[37] Crawford, Vincent P., and Nagore Iriberri. 2007a. "Level-k Auctions: Can a Non-Equilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?" Econometrica, 75(6): 1721:1770.

[38] Crawford, Vincent P., and Nagore Iriberri. 2007b. "Fatal Attraction: Salience, Naivete, and Sophistication in Experimental Hide-and-Seek Games." American Economic Review, 97(5): 1731-1750.

[39] De Giorgi, Enrico G. and Stefan Reimann. 2008. "The [alpha]-beauty contest: Choosing numbers, thinking intervals." Games and Economic Behavior, 64(2):470-486.

[40] Dyer, Douglas, John H. Kagel, and Dan Levin. 1989. "A Comparison of Naive and Experienced Bidders in Common Value Offer Auctions: A Laboratory Analysis." The Economic Journal, 99(394): 108-115.

[41] Georganas, Sotiris, Paul J. Healy, and Roberto A. Weber. 2010. "On the Persistence of Strategic Sophistication." Unpublished Ohio State working paper

[42] Gneezy, Uri. 2005. "Step-Level Reasoning and Bidding in Auctions," Management Science, 51(11): 1633-1642.

[43] Goeree, Jacob, Charles A. Holt, and Thomas R. Palfrey. 2002. "Quantal Response Equilibrium and Overbidding in Private-Value Auctions." Journal of Economic Theory. 104(1):247-272.

[44] Goeree, Jacob, Charles A. Holt, and Thomas R. Palfrey. 2005. "Regular Quantal Response Equilibrium." Experimental Economics, 8(4): 347-367.

[45] Goldfarb, Avi, and Botao Yang. 2009. "Are All Managers Created Equal?" Journal of Marketing Research, 46(5): 612-622.

[46] Goldfarb, Avi, and Mo Xiao. 2010. "Who thinks about the competition? Managerial Ability and Strategic Entry in US Local Telephone Markets." Unpublished.

[47] Hahn, P. Richard, Lum, Kristian, and Mela, Carl 2010. "A semiparametric model for assessing cognitive hierarchy theories of beauty contest games." Unpublished

[48] Haruvy, Ernan, and Dale O. Stahl. 2007. "Equilibrium selection and bounded rationality in symmetric normal-form games." JEBO 62(1):98-119.

[49] Haruvy, Ernan, Dale O. Stahl, D., and Paul W. Wilson. 1999. "Evidence for optimistic and pessimistic behavior in normal-form games." Economics Letters 63(3):255-259.

[50] Haruvy, Ernan, Dale O. Stahl, D., and Paul W. Wilson. 2001. "Modeling and testing for heterogeneity in observed strategic behavior." Review of Economics and Statistics 83(1):146-157.

[51] Ho, Teck-Hua, Colin Camerer, and Keith Weigelt. 1998. "Iterated Dominance and Iterated Best Response in Experimental 'p-Beauty Contests'." American Economic Review, 88(4): 947-969.

[52] Ivanov, Asen, Dan Levin, and Muriel Niederle. 2008. "Can Relaxation of Beliefs Rationalize the Winner's Curse?: An Experimental Study." Unpublished (August 27, 2008 version), available at http://www.hss.caltech.edu/ mshum/ec106/BeliefsExperiment.pdf.

[53] Ivanov, Asen, Dan Levin, and Muriel Niederle. 2010a. "Can Relaxation of Beliefs Rationalize the Winner's Curse?: An Experimental Study." Unpublished (April 5, 2010 version), available at
http://www.people.vcu.edu/ aivanov/20100405_CanBeliefsRationalizeWC_WP.pdf.

[54] Ivanov, Asen, Dan Levin, and Muriel Niederle. 2010b. "Can Relaxation of Beliefs Rationalize the Winner's Curse?: An Experimental Study." Econometrica, 78(4): 1435-1452.

[55] Jehiel, Philippe. 2005. "Analogy-Based Expectation Equilibrium." Journal of Economic Theory, 123(2): 81-104.

[56] Jehiel, Philippe, and Frederic Koessler. 2008. "Revisiting Games of Incomplete Information with Analogy-Based Expectations." Games and Economic Behavior, 62(2):533-557.

[57] Johnson, Eric, Colin F. Camerer, Sankar Sen, and Talia Rymon. 2002. "Detecting Failures of Backward Induction: Monitoring Information Search in Sequential Bargaining." Journal of Economic Theory, 104(1): 16-47.

[58] Kagel, John H., and Dan Levin. 1986. "The Winner's Curse and Public Information in Common Value Auctions." American Economic Review, 76(5): 894 - 920.

[59] Kagel, John H., and Dan Levin. 2002. *Common Value Auctions and the Winner's Curse.* Princeton University Press.

[60] Kawagoe, Toshiji, and Hirokazu Takizawa. 2008. "Level-K Analysis of Experimental Centipede Games." Unpublished

[61] Kawagoe, Toshiji, and Hirokazu Takizawa. 2009. "Equilibrium refinement vs. level-k analysis: An experimental study of cheap-talk games with private information." Games and Economic Behavior, 66(1), 238-255

[62] Kubler, D. and Weizsecker, G. 2004. "Limited depth of reasoning and failure of cascade formation in the laboratory." Review of Economic Studies, 71(2):425:441.

[63] Levin, Dan, John H. Kagel, and Jean-Francois Richard. 1996. "Effects and Information Processing in English Common Value Auctions." 86(3): 442-60.

[64] Malmendier, Ulrike, and Devin Shanthikumar. 2007. "Are Small Investors Naive about Incentives?" Journal of Financial Economics, 85(2), 457-489.

[65] Milosavljevic, Milica, Alec Smith, Christof Koch, and Colin Camerer. 2010. "Low-Level Visual Saliency Can Influence High-Level Cognitive Behavior: Evidence from Coordination Games." Caltech working paper.

[66] Nagel, Rosemarie. 1995. "Unraveling in Guessing Games: An Experimental Study." American Economic Review, 85(5): 1313-1326.

[67] Niederle, Muriel. 2010. "Intelligent Design: The Relationship of Economic Theory to Experiments: Treatment driven Experiments". Unpublished (available at http://www.stanford.edu/∼niederle/Intelligent%20Design.August15.pdf).

[68] Oestling, Robert, Joseph Tao-Yi Wang, Eileen Chou, and Colin Camerer. 2009. "Testing Game Theory in the Field: Swedish LUPI Lottery Games." Unpublished.

[69] Rey-Biel, Pedro. 2009. "Equilibrium Play and Best Response to (Stated) Beliefs in Normal Form Games." Games and Economic Behavior, 65(2):572-585.

[70] Rogers, Brian, Thomas R. Palfrey, and Colin Camerer. 2009. "Heterogeneous Quantal Response Equilibrium and Cognitive Hierarchies," Journal of Economic Theory, 144(4), 1440:1467.

[71] Rubinstein, Ariel. 2007. "Instinctive and Cognitive Reasoning: A Study of Response Times." The Economic Journal, 117(523):1243-1259.

[72] Sonsino, Doron, Ido Erev, and Sharon Gilat. 2002. "On Rationality, Learning, and Zero-Sum Bettings: An Experimental Study of the No-Betting Conjecture." Unpublished.

[73] Sovik, Ylva. 2009. "Strength of Dominance and Depths of Reasonings: An Experimental Study." Journal of Economic Behavior & Organization, 70(1-2): 196-205

[74] Stahl, Dale O.. 1996. "Boundedly rational rule learning in a guessing game." Games and Economic Behavior, 16, 303-330.

[75] Stahl, Dale O.. 1998. "Is step-j thinking an arbitrary modelling restriction or a fact of human nature?" Journal of Economic Behavior and Organization, 37, 33-51.

[76] Stahl, Dale O., and Ernan Haruvy. 2008. "Level-n bounded rationality and dominated strategies in normal-form games." Journal of Economic Behavior & Organization, 66(2):226-232.

[77] Stahl, Dale O., and Paul R. Wilson. 1994. "Experimental evidence on players' models of other players." Journal of Economic Behavior & Organization, 25, 309-327.

[78] Stahl, Dale O., and Paul R. Wilson. 1995. "On Players' Models of Other Players: Theory and Experimental Evidence." Games and Economic Behavior, 10(1): 218-254.

[79] Tor, Avishalom, and Max Bazerman. 2003. "Focusing Failures in Competitive Environments: Explaining Decision Errors in the Monty Hall Game, the Acquiring a Company Problem, and Multiparty Ultimatums." Journal of Behavioral Decision Making, 16: 353:374.

[80] Wang, Joseph Tao-Yi, Michael Spezio, and Colin F. Camerer. 2010. "Pinocchio's Pupil: Using Eyetracking and Pupil Dilation To Understand Truth:telling and Deception in Games." American Economic Review, 100(3): 984:1007.

[81] Weizsecker, Georg. 2003. "Ignoring the Rationality of Others: Evidence from Experimental Normal-Form Games." Games and Economic Behavior, 44(1): 145-171.

[82] Wright, James R., and Kevin Leyton-Brown. 2010. "Beyond Equilibrium: Predicting Human Behavior in Normal-Form Games." Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence.