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

Colin Camerer† Salvatore Nunnari‡ Thomas R. Palfrey†

December 11, 2014

Abstract

We report an experiment on a simple common value auction to investigate the extent to which bidding can be explained by quantal response equilibrium, in combination with different assumptions about the structure of bidder beliefs—the cursed equilibrium model and models that posit levels of strategic sophistication. Using a structural estimation approach, we find a close correspondence between the theoretical predictions of those models and experimental behavior. The basic pattern of average bids in the data consists of a combination of overbidding for low signals, and value-bidding for higher signals. The logit QRE model with heterogeneous bidders fits this pattern reasonably well. Combining quantal response with either cursed beliefs (CE-QRE) or a level-k of strategic sophistication (LK-QRE, CH-QRE) leads to a close match with the data. All these variations on quantal response models predict minimal differences of average bidding behavior across different versions of the game, consistent with the experimental findings. Finally, we reanalyze data from an earlier experiment on the same auction by Ivanov, Levin and Niederle (2010). While their data exhibit much more variance compared with ours, nonetheless, we still find that these models also fit their data reasonably well, even in the presence of extreme overbidding observed in that experiment. Overall, our study indicates that the winner curse phenomenon in this auction is plausibly attributable to limits on strategic thinking combined with quantal response.

*We are grateful for comments from seminar audiences at the 2010 North-American ESA Meeting in Tucson, the Security Market Auctions and IPOs Conference at Northwestern University, the University of Southern California, and Stanford University. We thank Asen Ivanov, Dan Levin, and Muriel Niederle for sharing their data and for detailed comments on an earlier draft. Palfrey thanks NSF and Russell Sage Foundation for financial support.

†Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, CA 91125.

‡Department of Political Science, Columbia University, New York, NY 10027

1 Introduction

In some markets, agents assign a common value to an object, which is equal to a particular statistic of their privately observed value signals. Bidding in such markets often leads to a "winner's curse", in which the winning bid is systematically greater than the common value.

Two recent studies of common-value auctions include an interesting and strategically simple case in which the common value is the *maximum* of the signals of two agents (Bulow and Klemperer, 2002). Carrillo and Palfrey (CP 2011) study this economic environment in the context of two-person bilateral trade. Ivanov, Levin and Niederle (ILN 2010) study a two-person second-price auction mechanism in which the common value is the maximum signal; we call this the maximum-value auction. In CP, the equilibrium prediction is that there should be no trade, regardless of which statistic determines the common value (including the maximum). In ILN, if the common value is the maximum signal, the unique symmetric equilibrium is to bid one's signal.

Both studies actually report large and significant deviations from the Nash equilibrium. That is, trade occurs frequently in the CP experiment, and bids in the ILN experiment are often much higher than the bidders' signals (in fact, a large fraction of bidders bid more than 10 times the maximum object value at least once). CP conjecture that the deviations they observed can plausibly be the result of bounded rationality of the form captured by cursed equilibrium (CE; Eyster and Rabin, 2005) or analogy-based expectations (Jehiel 2005). ILN present evidence that they interpret as discrediting explanations based on models of limited strategic sophistication, including both CE and level-k models (LK; Stahl 1996, 1998 Crawford and Irriberri 2007a). The latter conclusion is surprising, because LK models and the closely related cognitive hierarchy model (CH; Camerer, Ho, and Chong 2000) account for many features of observed behavior in a wide variety of other experimental games, including standard common value auctions, as well as measures of visual attention to payoffs and game structure[1] and brain activity[2].

Using the maximum-value auction as our starting point, this paper takes a dual approach to look deeper into that conclusion. First, we conduct a new set of laboratory auctions with the same rules and procedures of the primary treatments explored by ILN. This allows us to identify whether the surprising findings are the result of an unusual set of data. Second,

[1]Wang, Spezio and Camerer 2010, Brocas et al Forthcoming, Camerer and Smith 2012, Polonio et al 2014

[2]Coricelli and Nagel 2009; Bhatt et al. 2010, Yoshida et al 2010, Bhatt et al. 2012.

we adopt a structural estimation approach to evaluate the models of belief formation based on limited strategic sophistication and/or cursed beliefs. To provide an error structure for doing this, we combine these belief based models with quantal response equilibrium (QRE; McKelvey and Palfrey 1995).[3] The main finding is that bids are well accounted for by a quantal response equilibrium model of heterogeneous bidders, and this conclusion is even stronger with respect to the models that combine QRE with biased belief models of CE, LK, or CH.

The intuition for why these models do so well in this particular game is rather simple. While it is a fact that there is a unique symmetric Bayesian-Nash equilibrium in which bidders bid their signals, the equilibrium is weak; indeed it is just about as weak as an equilibrium can possibly be. If the other bidder is bidding according to the Nash equilibrium, i.e., bid=signal, then *every bid greater than or equal to your signal* is a (weak) best response to Nash equilibrium.[4] QRE imposes the assumption that beliefs about distribution of choices of other players are accurate, but allows for imperfect (noisy) response. In maximum-value auctions, a natural conjecture is that (symmetric) regular quantal response equilibria (Goeree, Holt and Palfrey 2005) will typically entail a substantial frequency of overbidding, even for rather small levels of noise, because bidding above one's signal is a small mistake (in fact, not a mistake at all in the Nash equilibrium). But then, if in equilibrium many bidders are overbidding substantially, underbidding is nearly a best response because the probability of winning when bidding one's signal is very low. The CE effects will further increase the amount of overbidding somewhat, as is also possible when beliefs are formed according to quantal response versions of LK or CH.

Following ILN, in our experiment subjects first participated in one treatment, which was followed immediately by a second surprise treatment. In the first treatment, *Phase I*, the signal distributions are uniform discrete distributions over the 11 integers 0-10, and are independent for the two bidders (randomly) paired together each round. In Phase I, subjects received each of the 11 integer signals exactly once – sampling without replacement– in a sequence of 11 rounds. Subjects could bid any (weakly positive) amount. In the second treatment, called Phase II, subjects bid against the distribution of *their own previous signal-bid pairs* sampled randomly from the Phase I auction against other people.[5] It is crucial to

[3]We limit the analysis to the maximum-value auction, ignoring CP, because the game is symmetric and the experimental implementation of the signal space is much coarser than the signal space in CP.

[4]Even underbidding can be optimal in the asymmetric Nash equilibria of the game.

[5]In ILN there are also two versions of Phase II, reminding subjects of their Phase I bids (*ShowBid*) or

note that to disable learning, there is no feedback throughout the entire sequence of 22 bids (=11 bids per phase $\times$ 2 phases).

Figure 1 shows scatter plots of actual bids and private signals in Phase I (panel a) and Phase II (panel b) of our experiments, excluding bids above \$15. There are two striking findings that are common to our data and ILN's. First, there is substantial overbidding which has a "hockey stick" shape: average bids when signals are below 5 are relatively flat, around $3-5$, only weakly responding to increases in the signal; bids when signals are above 5 are slightly above the signals, and steadily increase with signals. As we see below, this pattern arises very naturally from QRE. Second, average bids do not change very much between the Phase I two-player auction and the Phase II (bid-against-a-computerized-version-of your-own-phase-I) treatment. We show that these two empirical results are consistent with both quantal response equilibrium (QRE) and with quantal response versions of theories in which strategic thinking is limited by either cursed beliefs or levels of strategic sophistication.[6]

Note that accurate fits from QRE do not arise because of a loose intuition that 'quantal response can fit anything'. First, this assertion is simply not generally true for *regular* quantal response equilibria, including the most common logit and power specifications (Goeree, Holt and Palfrey, 2005). These equilibria satisfy various natural properties of continuity and responsiveness of choice probabilities to payoffs—higher expected payoffs lead to higher choice probabilities. If subjects' choices do not obey the latter property then regular QRE cannot fit the data well.

Second, our estimation procedure is designed to guard against overfitting the data. Specifically, the Phase I data are first used to derive QRE parameters, which fit those data best using maximum likelihood estimation. Then those estimates from Phase I are used to provide out-of-sample forecasts of Phase II behavior.[7] In principle, these out-of-sample forecasts for average bidding behavior should be fairly close to the Phase I average bids, because of the

not reminding them (*Baseline*). Our experiment uses only the ShowBid procedure to avoid concerns about subjects' memory.

[6]There is a third effect which is special to ILN's data: extreme overbidding. Figure 10 in the appendix reports the unconstrained bids by signal for Phase I, for both datasets. All bids are shown on a log scale (since bids can range from 0 to $1,000,000$). 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). Surprisingly, 27% of ILN's subjects bid 100 or more at least once. These surprisingly high bids were not observed in our data.

[7]Note that if the models were wrong, and were overfit to Phase I data, then applying those estimates to the Phase II could result in poor fits. The better the out-of-sample estimates fit the Phase II data, the more convincing the findings.

relatively small treatment effects in both our data and the ILN data. This is what we find.

1.1 The Winner's Curse

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 CE and the related ABE model. 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 strategic uncertainty, but there is identical conditional expectation. Indeed, there is some experimental evidence of such a mistake in decisions (see Carroll, Bazerman and Maury 1988's "beastie run"; Tor and Bazerman 2003; Charness and Levin 2009). This fact is especially relevant to the experiment at hand, because it implies that even when bidders are explicitly told the joint distribution of bids and signals of the opponent (as in Phase II bidding), they still will not fully take account of the correlation between their opponent's bids and action.

What about strategic models of the winner's curse? One approach is CE. Another class of models include the LK and CH specifications of beliefs which assume that individuals have different beliefs that correspond to different levels of strategic sophistication. Both specifications start with some clearly specified level-0 (L0) play. L0 choices are typically assumed to be random and uniform, but could more generally be considered instinctive (Rubinstein, 2007), focal, based on salience (Crawford and Iriberri, 2007a, 2007b; Milosavljevic, Smith, Koch and Camerer, 2010), or corresponding to other types (Stahl and Wilson, 1994, 1995).

Level-1 (L1) players assume they are playing L0 players.[8] An iterated hierarchy then follows, most typically in two different ways: in "level-k" models Lk players (for $k \geq 1$) believe

[8]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).

their opponents are all using Lk-1 reasoning. In CH models Lk 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 exact best responses given their beliefs. However, the earliest forms of level-k and CH models typically replace this best response assumption with some kind of quantal response (Stahl and Wilson 1994, 1995; Ho, Camerer and Weigelt 1998; Costa-Gomes, Crawford and Broseta 2001, Chong, Ho and Camerer 2005), or an interval of responses near a single best response (Nagel 1995).[9] Indeed, models combining quantal response with LK or CH account for 60% of papers (37 out of 61) using level-k or CH.[10]

This observation raises two interesting empirical questions: First, can quantal response combined with CE, LK or CH beliefs (denoted CE-QRE, LK-QRE, or CH-QRE) or equilibrium beliefs (QRE)[11] explain the bids in maximum-value experiments? And second, can such models explain why bids do not change much in the Phase II bid-against-yourself treatment? We answer these questions by estimating those models on Phase I data, seeing how well they fit, and then testing how well they fit Phase II bids (out of sample).

1.2 The Maximum Value Game

In the maximum-value second price auction, two players observe private signals x_1 and x_2 drawn independently from a commonly known distribution over the real line. 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. 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

[9] 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.

[10] See Camerer et al. 2012 for more details.

[11] Goeree, Holt and Palfrey (2002, 2005).

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 symmetric equilibrium where both players bid their signal can be solved for with two rounds of elimination of weakly dominated strategies, as shown in ILN.

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 (and hence are not necessarily mistakes). 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 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 from laboratory experiments.

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

This section explains the three key ideas about the incentives to deviate from equilibrium behavior. First, we explain some intuitions of how bidding in Phase I depends on the change from an assumption of *perfect* best response behavior to *quantal* response behavior. Second, we explain the intuition for the effect of imperfect beliefs, à la CE and LK[12] models, when combined with quantal response behavior. Third, we provide a theoretical intuition for why the difference in bidding behavior between Phase I and Phase II can be negligible under both the quantal 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

[12]For most of the analysis of models of strategic sophistication that follows we focus on LK rather than CH, because it is easier to explain. The implications of those two models in this auction game are similar. We fit the data using both models and obtain nearly identical results.

opponent, under different hypotheses about his or her degree of quantal response behavior. Recall that in the logit version of quantal response, for each possible signal x, a bidder uses a behavior 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, for each $b \in B$ and for each signal x:

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

where B is the set of available bids, and $U(b|x)$ denotes the expected utility of bidding b conditional on privately observing signal 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, $U(b|x)$. QRE imposes rational expectations, or correct equilibrium beliefs, so $U(b|x)$ is computed using the QRE bidding strategies for a given value of λ.

At one extreme, if λ equals infinity, then everyone believes that everyone is bidding perfectly, in the sense of bidding according to the unique symmetric Nash equilibrium. Thus beliefs boil down to believing that one's opponent will always bid his signal. However, if this is the case, then overbidding any amount is also a perfect best response, 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 $b = x$. In contrast, there are significant losses from underbidding if one's opponent is bidding his signal. This asymmetry of payoff losses has important consequences for QRE bidding strategies.

Next, consider the opposite extreme, λ equal to 0, where everyone believes everyone is bidding completely randomly. 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 is simply $E[\max\{x_i, x_{-i}\}|x_i] = \frac{x_{you}(x_{you}+1)}{11} + \frac{(10-x_{you})(11+x_{you})}{22}$. Straightforward calculations show that your optimal bid when $x_i = 0$ equals 5 and the optimal bid increases monotonically to 10 when your value is 10. For all values of x_i except 10, the optimal bid ($E[\max\{x_i, x_{-i}\}|x_i]$) is strictly greater than x_i. Hence, the direct effect of quantal response behavior involves a combination of over-and underbidding, while the

Figure 1: Bids as a function of private signal values in our experiment: Phase I (a), and Phase II (b). Bids above 15 are not shown. The continuous line is the median bid as a function of private signal (considering all bids).

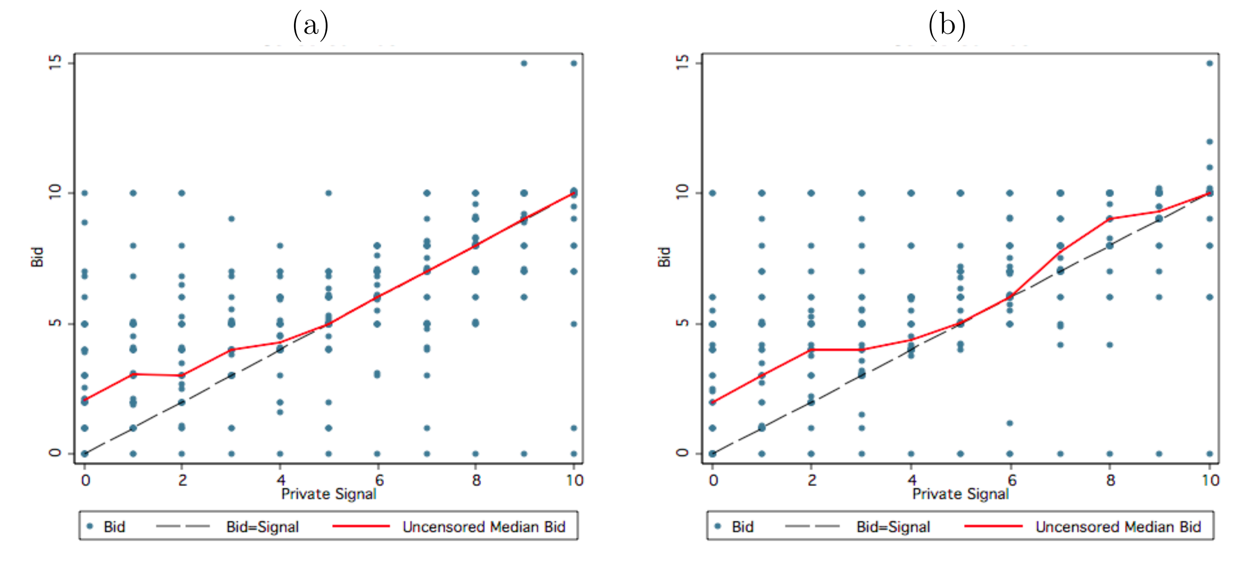

Figure 2: QRE strategies for different λ's, Phase I

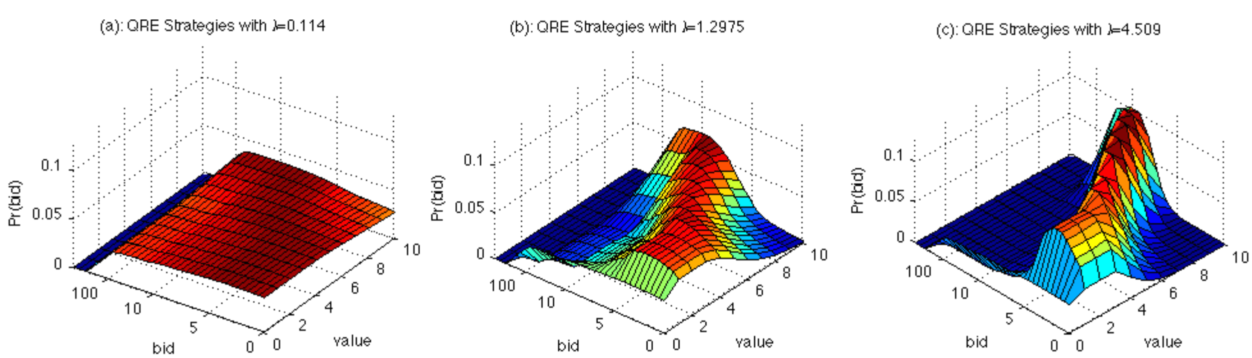

equilibrium effect of quantal responses pushes in the direction of overbidding.

Figure 2 shows three histograms of the theoretical joint distribution of bids and signals in the logit QRE, for three values of λ (0.114, 1.2975, 4.509). The distribution begins rather flat for low values of λ (left panel). As λ increases (middle and right panel) the distribution of bids becomes more concentrated, with substantial overbidding for all signals. 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 exact joint distribution of bids and signals 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, Palfrey, and Camerer (2009).

The effects of CE and LK complement, and in some cases reinforce, the effects of quantal response equilibrium. With respect to CE, a stark way to see this is to note that if $\lambda = 0$ there is no winner's curse, so quantal response equilibrium, for *any* positive value of λ will necessarily embody 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 LK models, L1 bidders behave exactly as if they are facing $\lambda = 0$, so the discussion of optimal responses is the same as described above. However, L2 bidders who think they are playing L1 know they are playing overbidders so *an* optimal response is to bid their signal. In fact L2 bidders generally have a range of optimal bids conditional on signal. To consider even higher levels, recall that if an opponent is bidding their signal or lower, then overbidding and bidding one's signal are both weak best-responses. Therefore, an L3 bidder is actually expected to overbid against an L2 bidder. Following a similar logic, higher level players alternate from level to level between over and underbidding (or value bidding) with a range of optimal bids for each signal (k even) and (approximately) signal bidding except for the lowest signals (k odd). Table 1 gives the range of optimal bids for levels, $k = 1, 2, 3, 4$. Notice that all four level types may overbid with low signals, and levels 0, 2, and 4, may overbid or underbid, even with the same signal, and million dollar bids are best responses for even types for some signal realizations.[13]

Table 1: Optimal Bids in Level-k Model

Signal	*Level*			
	1	2	3	4
0	5.00	$0.00 - 4.99$	2.50	$0.00 - 2.49$
1	5.09	$0.00 - 4.99$	2.67	$0.00 - 2.49$
2	5.27	$0.00 - 4.99$	3.00	$0.00 - 2.49$
3	5.55	$0.00 - 4.99$	3.50	$2.68 - 3.49$
4	5.91	$0.00 - 4.99$	4.17	$3.51 - 4.16$
5	6.36	$0.00 - 5.00$	5.00	$5.00 - 10^6$
6	6.91	$5.92 - 6.35$	6.00	$6.00 - 10^6$
7	7.55	$6.92 - 7.54$	7.00	$7.00 - 10^6$
8	8.27	$7.56 - 8.26$	8.00	$8.00 - 10^6$
9	9.09	$8.28 - 9.08$	9.00	$9.00 - 10^6$
10	10.00	$9.10 - 10^6$	10.00	$10.00 - 10^6$

Thus, the LK model assuming best response, predicts a mixture of random bidding, overbidding, and signal-bidding, as in QRE.[14]

What are the implications of these models of quantal response and imperfect beliefs for Phase II bidding behavior? In Phase II, each bidder competes against a computerized player that is programmed to bid exactly according to the bidding function the bidder used in Phase I. Because there is no feedback about payoffs and auction outcomes between phases, in our analysis we assume that the degree of payoff responsiveness or cursedness is the same in Phase I and Phase II for each bidder, but each bidder still takes fully into account that their opponent's (stochastic) bidding strategy is governed by a different distribution in Phase I and Phase II.

First, consider the QRE model. In this case, if a bidder is a low-λ 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 also reflected in the data. By the same token, bidders who were bidding

[13]For k sufficiently large, odd types are value bidders and million dollar bids are best responses for even types for all signal realizations.

[14]Optimal bids for L2 and greater in the CH model are somewhat different from the LK bids shown in the table, and depend on the relative frequencies of lower levels.

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 across the two phases.

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 that a human bidder who fell prey to cursed beliefs in Phase I would have an epiphany between phases 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 otherwise 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 LK models for Phase II. With a *perfect* best response model, a LK bidder in Phase I should behave like a Lk+1 bidder in Phase II. For example, if a bidder is a value bidder (k odd) in Phase I she should be an overbidder in Phase II and viceversa. This is not what one observes in the data. However, the picture changes in two interesting ways if there is one considers quantal response behavior instead of perfect best responses. First, bidders who are L0 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 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.

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 L1'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 bidding behavior in the comparison of Phase I to Phase II.[15]

[15]The same conclusion holds for the CH model with quantal response, which predicts identical L1 behavior

2 Procedures and Data

Signals are uniformly, independently, and discretely distributed from 0 to 10 in integer increments. The baseline auctions are called "Phase I" auctions, and strategies are elicited using a particular form of the "strategy method".[16] This is done by eliciting from each bidder one bid for each of the 11 possible signals in a random sequence. Bids could be any non-negative amount up to $1,000,000.00, in penny increments. Payoffs are determined for each of these signal-bid pairs by random matching with an opponent, but there is no feedback until both phases of the experiment finish. Payoffs are converted to U.S. dollars using an exchange rate of $0.50US per experimental dollar. Subjects were paid $20 plus or minus the total of all the payoffs from all auctions. Notice that as is the case in most common value experiments it is possible to go bankrupt, if you overbid persistently. In our experiment, no subject lost money from the auctions.Phase II auctions were conducted immediately following the Phase I auctions. In Phase II each subject again submitted 11 bids, one for each possible signal. These bids were then randomly matched against a computer that was programmed to use exactly the same bidding function that had been used by the subject in Phase I. Subjects were told this before preparing their Phase II bids.[17].

3 Model Estimation and Forecasting

The figures presented in the Introduction, while informative of the overall pattern of bidding, conceal the extensive heterogeneity in bidding behavior across subjects. ILN recognize the importance of heterogeneity and report their results in the form of a heuristic classification of bidding strategies. A more extensive classification analysis is contained in Costa-Gomes and Shimoji (2011). To illustrate the degree of heterogeneity in a very compact graphical format Figure 3 graphs the empirical inverse cumulative distribution function of (unconditional) average bids by subject for Phase I data. There is a significant degree of heterogeneity: 24% of subjects have an average bid smaller than 5; 35% between 5 and 6; 41% greater than 6.[18]

Because of the importance of heterogeneity, we consider in this section three different models of bidding behavior that incorporate different kinds of heterogeneity. The first kind

---

and a similar L2 response to the Phase II treatment.

[16]See Selten (1967) and Roth (1995, p. 322-323), for example.

[17]Our procedures were identical to a subset of ILN procedures, including instructions and software

[18]The average bid of a value bidder is 5.

Figure 3: CDF of Mean bids, Phase I

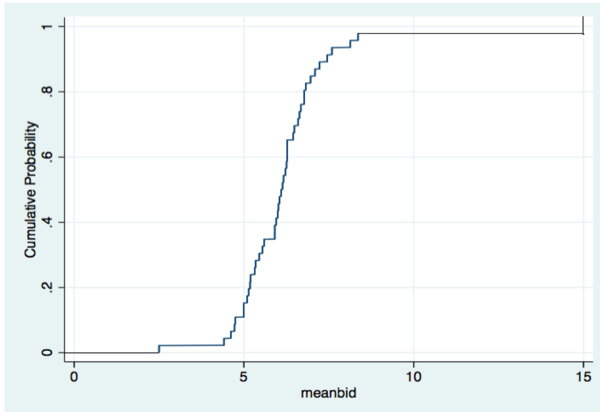

of heterogeneity we consider is the degree to which subjects are best-responding. This is *heterogeneity with respect to choice* (specifically the response parameter λ), but maintaining the assumption that all players share homogeneous beliefs about the strategies of the other players. Significant evidence of such heterogeneity has been reported elsewhere and can be thought of as a reduced form approach to modeling differences in skill or carelessness of the players. For reasons discussed earlier, the relatively flat payoffs and the fact that best responses are generally not unique in this game suggest that this kind of heterogeneity may help explain some of the variation of bidding behavior across players.

The second kind of heterogeneity, captured by the LK and CH modeling approach, has also been discussed earlier and explicitly incorporates a model of *heterogeneity with respect to beliefs*; it is one of the primary motivations behind the experimental design. In our estimation, we assume that L players believe all their opponents are (quantal responding) Lk-1 players, and that the frequency of Lk players in the population follows a Poisson distribution.[19]

[19] In addition to this, we estimate the CH version of the level-k model (Camerer, Ho, and Chong 2004) where step k players believe 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 other players are step 0 players, step 2 players assume the other players are a combination of step 0 players and (quantal responding) step 1 players, and so on. The estimates for this alternative model are very similar to the estimates for the regular version of level-k.

3.1 QRE and HQRE Predictions

In the logit version of HQRE all players have common rational expectations (and hence no belief heterogeneity) as in QRE, but can have different response parameters.
That is, players exhibit *choice heterogeneity.* Thus each player uses a behavioral strategy where the log probability of choosing each available action is proportional to its (rational expectations) expected payoff, where the proportionality factor, λ_i varies across subjects. In particular, for our setting we have:

$$Pr(b|x) = \frac{exp(\lambda_i U(b|x))}{\sum_{a\in B} exp(\lambda_i U(a|x))} \tag{1}$$

for all $b \in B$ and for each signal x.

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

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

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). Given these expected payoffs, the choice probabilities follow from the choice rule in (1), i.e.:

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

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 auctions to compute the empirical expected utility of each bid given each private signal, $\widehat{EU}(b_i|s_i)$. 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. This is how we compute $\widehat{EU}(b_i|s_i)$ for each (binned) bid b_i and each signal s_i in Phase I:

- $Pr(b_i > b_j)$ is the fraction of bids submitted by all subjects in Phase I lower than b_i;
- $E(max(s_i, s_j)|b_i > b_j)$ is the average value of the object in all Phase I auctions where a subject bids less than b_i;
- $E(b_j|b_i > b_j)$ is the average bid among all bids strictly lower than b_i submitted by all subjects in Phase I;
- $Pr(b_i = b_j)$ is the fraction of bids submitted by all subjects in Phase I equal to b_i;
- $E(max(s_i, s_j)|b_i = b_j)$ is the average value of the object in all Phase I auctions where a subject bids exactly b_i.

Then, for each subject i, we estimate, using MLE, the response parameter $\widehat{\lambda}_i$ that best fits i's 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 Phase I. Parameter estimates are summarized in Table 2.

Figure 4: HQRE Predictions

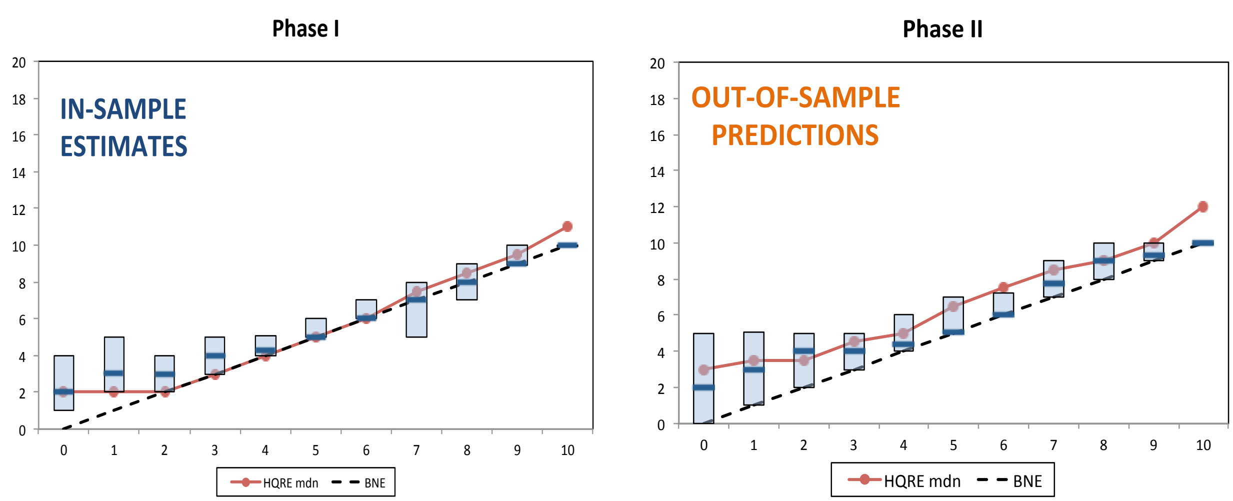

Table 2: Estimated Parameters

# Individuals	46
# Observations	506
QRE	
$\widehat{\lambda}$	3.59
-Log-likelihood (P1)	1447.8
HQRE	
$\widehat{\lambda}_i$ Quartiles	(3.40, 5.12, 8.70)
-Log-likelihood (P1)	1317.6
CE-QRE	
$\widehat{\chi}$	0.1
$\widehat{\lambda}$	3.66
-Log-likelihood (P1)	1440.4
CE-HQRE	
$\widehat{\chi}$	0.2
$\widehat{\lambda}_i$ Quartiles	(3.05, 6.55, 11.81)
-Log-likelihood (P1)	1289.7
LK-QRE (4 Types)	
$\widehat{\tau}$	2.7
$\widehat{\lambda}$	12.93
Types Frequency	(9%, 25%, 34%, 31%)
-Log-likelihood (P1)	1346.3
LK-HQRE (4 Types)	
$\widehat{\tau}$	4
$\widehat{\lambda}_i$ Quartiles	(3.82, 11.06, 47.51)
Types Frequency	(4%, 17%, 34%, 45%)
-Log-likelihood (P1)	1280.9

Figure 5: QRE Predictions

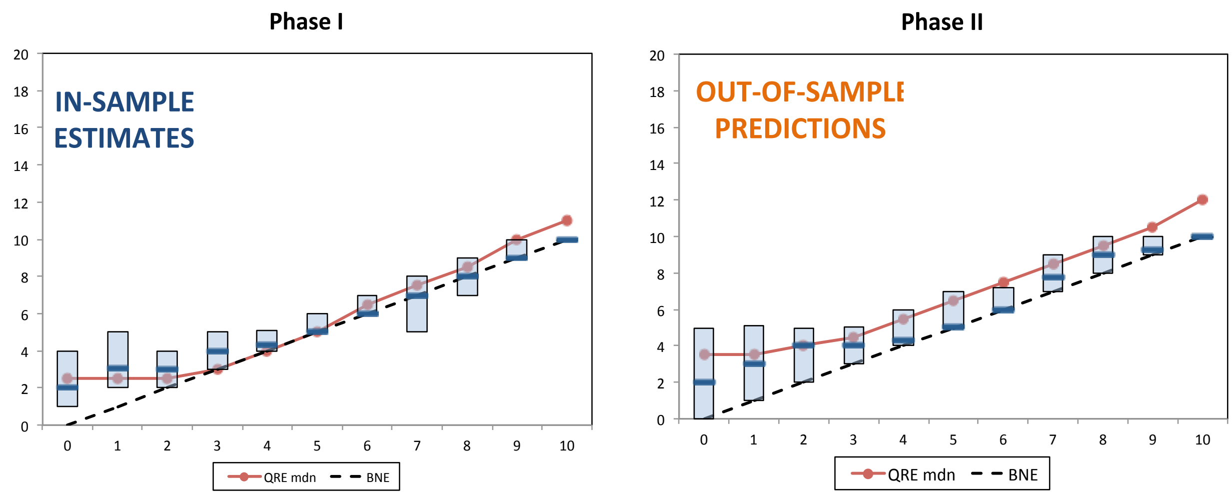

In the left panel of Figure 4 (for a model with individual specific λ's) and Figure 5 (for a model with an homogenous λ for all subjects), we show the quartiles of this distribution and we compare them with the quartiles of the observed bids (represented with boxplots). We emphasize that the reported predictions for Phase II (right panels of Figure 4 and 5) are, *out-of-sample predictions based on parameters estimated from Phase I.* For each subject, we first compute $\widehat{EU}_i(b_i|s_i)$ using the empirical joint distribution of bids and signals from his previous play in Phase I (i.e. each agent plays against his exact previous bidding function). This is how we compute $\widehat{EU}_i(b_i|s_i)$ for each (binned) bid b_i, each signal s_i, and each subject i in Phase II:

- $Pr(b_i > b_j)$ is the fraction of bids submitted by i in Phase I lower than b_i.
- $E(max(s_i, s_j)|b_i > b_j)$ is the average value of the object in all Phase I auctions where i bids less than b_i;
- $E(b_j|b_i > b_j)$ is the average bid among all bids strictly lower than b_i submitted by i in Phase I;
- $Pr(b_i = b_j)$ is the fraction of bids submitted by i in Phase I equal to b_i.
- $E(max(s_i, s_j)|b_i = b_j)$ is the average value of the object in all Phase I auctions where i bids exactly b_i.

Then, we predict his behavior in Phase II by plugging in these expected payoffs and the $\widehat{\lambda}_i$ estimated for Phase I into equation (2).[20]

3.2 CE-QRE and CE-HQRE Predictions

An agent i who is *cursed*, in the sense of Eyster and Rabin (2005), has correct expectations about the marginal distribution of other players' bids, but has incorrect beliefs about the joint distribution of bids and signals; specifically, a (fully) cursed bidder fails to account for the fact that bids are correlated with private signals. As a result 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 actual 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:

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

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

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

For estimation, as in Section 3.1, we use the empirical joint distributions of signals and bids in Phase I auctions to compute emprical values of $\widehat{EU}^{\chi}_{CE}(b_i|s_i)$. $\widehat{EU}(b_i|s_i)$ and most elements in $\widehat{EU}_{CE}(b_i|s_i)$ are computed as described in the previous section. One element in $EU_{CE}(b_i|s_i)$ is not part of $EU(b_i|s_i)$, that is, $E(max(s_i, s_j))$. This is the unconditional value

[20] 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.

of the object given your signal, that is, the value of the object assuming there is no correlation between the opponent's signal and his bid. As explained in Section 1.3, we compute this as follows:

$$E[\max\{s_i, sj\}|s_i] = \frac{s_i(s_i+1)}{11} + \frac{(10-s_i)(11+s_i)}{22}$$

Then, we estimate with a standard MLE procedure the vector of individual $\widehat{\lambda}_i$'s and the cursed parameter $\widehat{\chi}$ that best fit the data. 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 exact distribution of each bidder's own previous bids in Phase I) and, thus, different expected payoffs. The predicted distributions for Phase I are shown in the left panel of Figure 6 (for a model with individual specific λ's) and Figure 7 (for a model with an homogenous λ for all subjects). The out-of-sample predictions for Phase II are in the right panel of the same figures (using the estimated χ and λ's values from Phase I).

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.

Figure 6: CE-HQRE Predictions

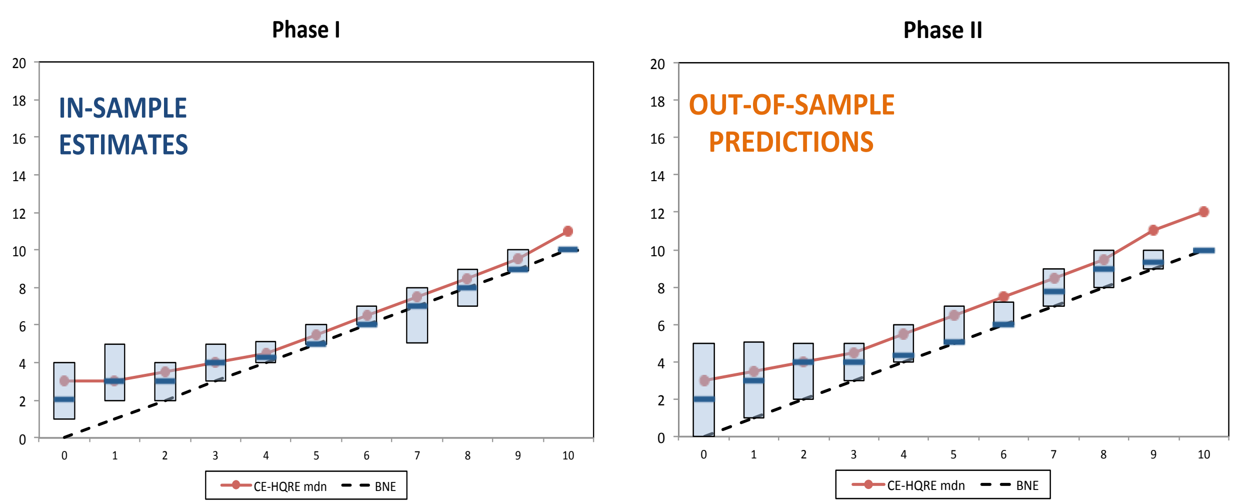

Figure 7: CE-QRE Predictions

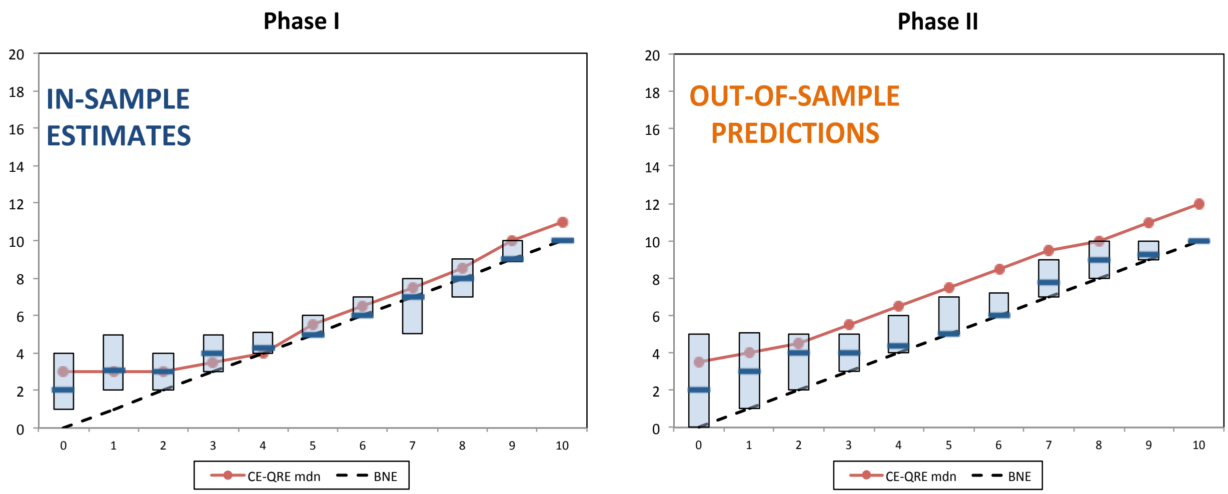

3.3 LK-QRE Predictions

For the LK estimation, denote the expected utility of step 1 players as $EU_{LK_1}(b_i|s_i, \tau)$, the one of step 2 players as $EU_{LK_2}(b_i|s_i, \tau)$, and the one of step 3 players as $EU_{LK_3}(b_i|s_i, \tau)$. We assume the presence of up to 4 levels of beliefs distributed according to a Poisson frequency distribution $f(k)$ of step k players in the population, with mean level equal to τ.[21] Furthermore bidders respond imperfectly as in a QRE with a homogeneous λ:

$$Pr_k(b_i|s_i) = \frac{exp(\lambda EU_{LK_k}(b_i|s_i))}{\sum\limits_{a \in B} exp(\lambda EU_{LK_k}(a|s_i))} \tag{3}$$

For each pair (λ, τ), (3) implies a joint distribution of bids and signals for type 1, type 2, and type 3 players (and, combining this with the random behavior of type 0 players, an aggregate distribution of bids). To generate the predictions in the left panel of Figure 8, we estimate with MLE the pair $(\widehat{\lambda}, \widehat{\tau})$ that fits best the empirical strategies from Phase I. Contrary to what we do for QRE and CE-QRE, we do not use the empirical joint distributions of signals and bids (i.e. the distributions observed in the experiments) in Phase I auctions to compute

[21] That is, the proportion of Lk in the population is $f(k) = e^{-\tau}\tau^k/k!$, which is characterized by one parameter $\tau > 0$. To keep the estimation simple, we only consider types $k \leq 3$. Allowing for one extra level in the estimation does not improve the fit.

the empirical expected utility of each bid given each private signal. Rather, we compute the expected utility of each bid given each private signal and the theoretical behavior of the lower level types. The expected utility of L1's is computed assuming they are facing L0's who choose any bid with the same probability, regardless of the signal. The expected utility of L2's (L3's) is computed assuming they are facing L1's (L2's) who imperfectly respond to L0's (imperfectly responding L1's), as in a QRE with a homogeneous λ. As for the other models, we bin bids for the estimation, for tractability. 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 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. The out-of-sample predictions for Phase II of the same experiments are shown in the right panel of Figure 8.

Figure 8: LK-QRE Predictions

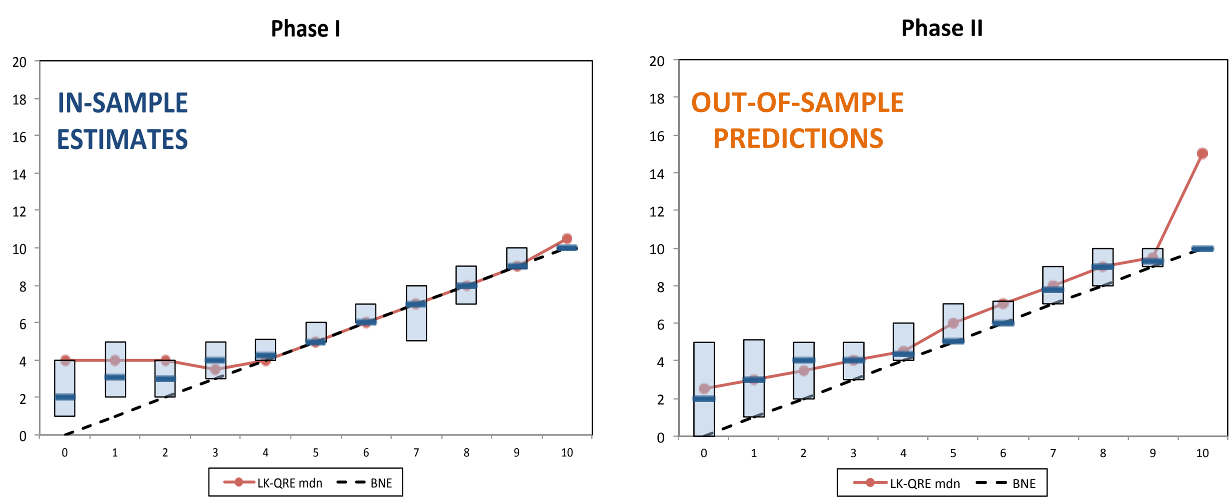

3.4 LK-HQRE Predictions

Finally, we estimate a level-k model with heterogeneous λ's, as we have done for the other models. As before, we assume the presence of up to 4 levels of beliefs, distributed according to a Poisson frequency distribution $f(k)$ of step k players in the population, with mean level equal to τ. We assume that each agent is randomly drawn from this distribution and that he responds imperfectly as in an HQRE with an individual specific λ_i:

$$Pr(b_i|s_i) = f(0)\frac{exp(\lambda_i EU_{LK_0}(b_i|s_i))}{\sum_{a\in B} exp(\lambda_i EU_{LK_0}(a|s_i))} + f(1)\frac{exp(\lambda_i EU_{LK_1}(b_i|s_i))}{\sum_{a\in B} exp(\lambda_i EU_{LK_1}(a|s_i))} + f(2)\frac{exp(\lambda_i EU_{LK_2}(b_i|s_i))}{\sum_{a\in B} exp(\lambda_i EU_{LK_2}(a|s_i))} + f(3)\frac{exp(\lambda_i EU_{LK_3}(b_i|s_i))}{\sum_{a\in B} exp(\lambda_i EU_{LK_3}(a|s_i))} \tag{4}$$

where $f(k) = e^{-\tau}\tau^k/k!$ is the probability a subject belongs to level k.

For each pair (λ_i, τ), (4) implies a joint distribution of bids and signals for a subject. To generate the predictions in the left panel of Figure 9, we estimate with MLE the vector ($\widehat{\lambda}_1$, $\widehat{\lambda}_2$,...,$\widehat{\lambda}_N$, $\widehat{\tau}$) that fits best the empirical strategies from Phase I (where N is the number of subjects in the experiment). For the out-of-sample predictions for Phase II we use the same $\widehat{\lambda}_i$ estimated for Phase I and compute the corresponding HQRE 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 lower level players. In addition to the fact that $\widehat{\lambda}_i$ is potentially different for every agent, we also 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. The out-of-sample predictions for Phase II of the same experiments are shown in the right panel of Figure 9.

3.5 Estimation and Forecasting on ILN Data

We perform the same analysis on data from ILN. Their data exhibit much more variance compared with ours and, for this reason, the parameters estimated for their data are quantitatively different than the parameters estimated for ours. Nonetheless, we still find that these models combining quantal response and limits on strategic thinking fit their data reasonably well, even in the presence of extreme overbidding observed in that experiment. The details are presented in the Appendix.

Figure 9: LK-HQRE Predictions

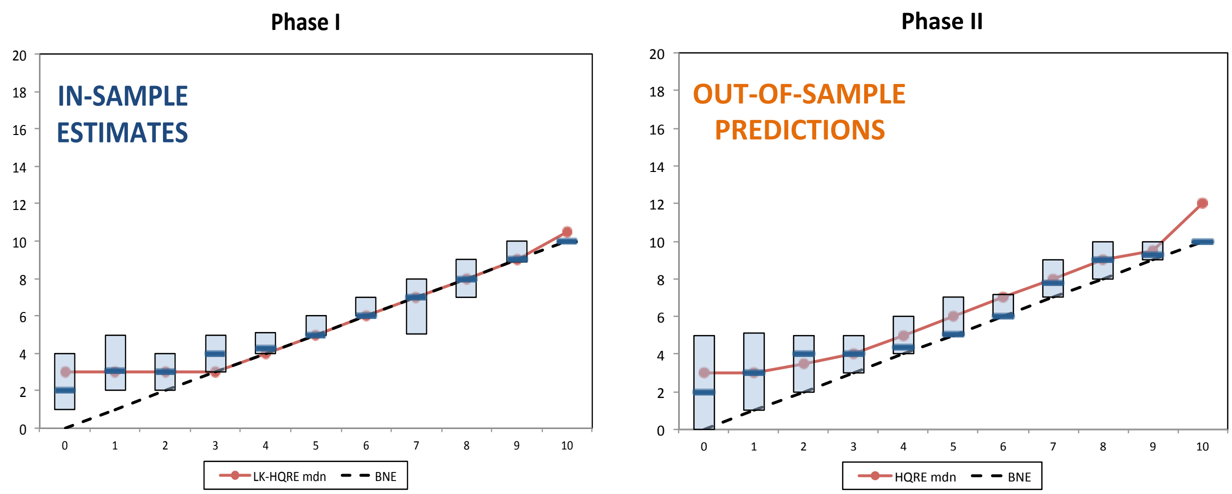

4 Play-Against-Yourself Phase II in P-Beauty Contests

The methodological technique of having people play against their own previous decisions is a novel idea introduced by ILN. First we comment on the unusual philosophy underlying this choice (see also the discussion in Crawford, Costa-Gomes and Iriberri 2013). Nonequilibrium models of beliefs (CE, ABEE, CH and level-k) are joint specifications of particular beliefs and response functions for different types of players. In principle, one could try to test the specification further by measuring subjects' beliefs in some other way to see if they appear to be consistent with (perceived) level-0 play.[22]

Asking a level-1 player to play their own previous choices *does not* test these nonequilibrium models because it does not test the joint specification of beliefs and choice. Instead, it puts the player into an entirely single-person new decision problem in which beliefs are exogeneously controlled by the experimenter. The nonequilibrium models are not meant to apply to this change, and hence make no prediction about what will happen.[23] This is why

[22]This is essentially what eyetracking, mousetracking, fMRI and other physical measurements try to accomplish, by seeing whether information or neural circuitry that are necessary to compute a belief are acquired or used.

[23]Niederle (in press) describes the logic of such an "elimination design": "Change the environment in a way the problem mostly stays the same, while, however, eliminating the conditions that allow the theory at hand to account for the phenomenon.... If behavior remains unchanged, at least it implies that other factors may be at work as well." That is, suppose a theory T predicts that in environment E, data D will result in

we, and many others working with nonequilibrium models, do not think the ILN Phase II design implies anything fundamental about such models.

Nonetheless, it is an empirical fact that the bid distributions do not change much between phases I and II in the maximum value game. Why not? The empirical success in predicting Phase II from Phase I estimation by QR models of all types suggests one answer. Another possibility was raised by one of the ILN authors, who wrote (Niederle, in press): "Though maybe, in this game, participants are so much at a loss, that the [level-k] model simply does not apply, that maybe, using it to explain behavior in common value auctions was simply too ambitious a goal?"

With this rough concept of participants 'at a loss' in mind, useful clarification may come from using the Phase II play-against-yourself method in a simpler game (with no private information). We therefore conducted a p-beauty contest game in two experimental sessions (for a total of 38 participants) at UCLA. Subjects were instructed to play a standard game choosing a number from [0, 100]. They were grouped with two others and whomever was closest to (2/3) of the three-person average would win $5. In a surprise Phase II, they were then told to copy their own previous number choice $X+1$, and their own number $X-1$, and use those numbers as choices by two artificial opponents. Then they chose a second Phase II number in a three-"person" contest against those perturbed previous choices.

An optimizing player will choose $(2/3)X$. Any choice $Y < X-1$ will win for $X < 3$ and $Y \in (X-1, X+1)$ will win for $X < 3$. Figure 10 plots the Phase I and Phase II bids for each subject. The distribution of Phase I choices looks similar to those from many other studies (see Camerer, 2003; Bosch-Domenech et al. 2002).[24]. The nearly-optimal shift in number choices from Phase I to II is readily apparent: Almost all subjects (80%) chose a winning number in Phase II, and the regression slope is .741 (standard error=.099), insignificantly different from the optimal slope of .667. Those choosing around 33 would typically be classified around level-1 (a more nuanced QRE estimation notwithstanding); as is evident, virtually all of them lower their choices in Phase II.

The Phase II results in this simple game are quite different from those in the maximum-value games. The difference undoubtedly has something to do with simplicity in computing

(T(E)=D). Now suppose you also observe D in environment E' in which T does not apply. The idea is that this suggests the possibility of a unifying theory T' in which T'(E)=T'(E')=D. Note, however, that if theory T has already explained data D_x well in many environments E_x (i.e. $T(E_x) = D_x$), the rival unborn theory T' must also backtrack and show that $T'(E_x) = D_x \forall x$ to eventually replace T.

[24]The summary statistics are: mean 21.89, median 16, standard deviation 19.01

Figure 10: Phase II vs Phase I Number Choices in P-Beauty Contest Game

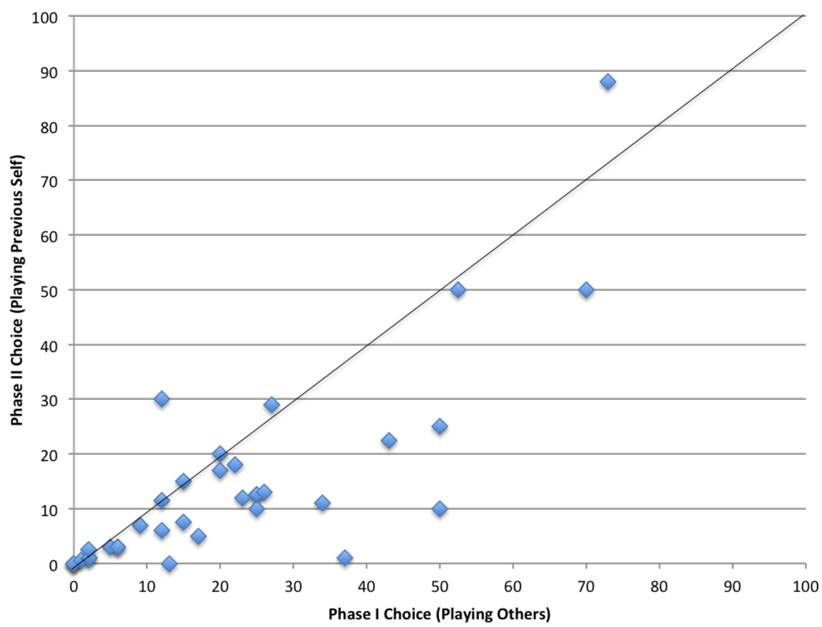

how to respond to one's own previous choices. Exploring precisely how that simplicity can be understood is a project for future work.

5 Discussion and Conclusion

Bidding behavior in the maximum value auction game exhibits a hockey stick-shaped pattern of overbidding. With low signals there is significant overbidding, and these overbids are relatively flat with respect to signals. With higher signals, overbidding is much less, with median bids closely tracking signals. The heterogeneous quantal response equilibrium model fits this pattern in the data, based on both within sample estimation and out-of-sample prediction. That is, that model clearly captures the hockey-stick bidding and it also generates approximately the right degree of variation in bids, measured by interquartile range. The fit is significantly and qualitatively improved by weakening the assumption of equilibrium beliefs (in QRE) with single parameter versions of cursed equilibrium or the level-k model with quantal response.[25] The addition of those non-equilibrium CE or level-k beliefs raises the level of estimated bids for lower signals, and lowers the level of out of sample predicted bids with high signals, to closely match the observed pattern of average bidding behavior.

It is important to emphasize that the predicted median bids in the Phase II data, which closely track the empirical median bids, are based on out-of-sample estimates obtained from the Phase I data. Moreover, in the *quantal response* models we explore here, the bidding behavior in Phase II is predicted to be quite similar to bidding behavior in Phase I. This is in stark contrast to models of *best response* behavior, where sharp differences in bidding behavior should be observed between Phase I and Phase II.

The distinction between best reply models and quantal response models is an important one because it has implications about how one evaluates the performance of belief based models based on either rational expectations (Nash equilibrium and QRE) or systematically incorrect beliefs (cursed equilibrium, ABEE, level-k, and cognitive hierarchy). If one insists on perfect best reply behavior, then all of the various belief-based models fail badly to explain the data from maximum value auction experiments. In contrast, if one allows for stochastic choice in the form of quantal responses, then a careful analysis of the data leads to the opposite conclusion, for a range of different assumptions about the structure of bidder beliefs.

We conclude with two general remarks that broadly relate to the findings of this paper. First, the maximum-value auction and associated treatments are not that useful for comparing different theories with equilibrium or non-equilibrium beliefs, under best response. All competing theories are soundly rejected by the Phase I behavior alone, in both the ILN data and the data reported here. However, in contrast, the game is actually very useful for the best response and quantal response approaches. Best response models cannot explain the distribution of bids and lack of treatment effects, while quantal response models can and do. Furthermore, the quantal response approach also allows one to compare belief based models that assume rational expectations with models of non-equilibrium beliefs, and to measure the degree to which non-equilibrium beliefs are an important behavioral factor. It turns out that non-equilibrium beliefs are an equally important factor in both experiments, in spite of a rather dramatic cosmetic difference in the two data sets. In one (ILN), there is a huge amount of variance, voracious overbidding, and a great deal more heterogeneity across the bidders; in the other (CNP) only 12 out of 1012 bids exceed the maximum value of 10. The

[25]The cognitive hierarchy model with quantal response fits the data similarly to the level-k model.

quantal response approach allows one to rigorously separate out the “noise” component of behavior from the systematic behavioral factors, which is virtually impossible if one views the data through the lense of perfect best-responses.

Second, 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 (Carroll, Bazerman, and Maury 1988; Charness and Levin, 2009), common value bargaining settings (Carrillo and Palfrey, 2009, 2011), and other common value games (Esponda and Vespa, 2014) clearly point to a widespread failure to correctly compute—or perhaps even comprehend—conditional expectations. However, by their very construction, models specifically constructed to model the difficulty of computing conditional expectations cannot possibly explain deviations from equilibrium in complete information games or in Bayesian games with a private values structure.[26] An advantage of the QRE approach, as well as the level-k and CH models, is that they have bite in nearly all games, that include but are not limited to only common value environments. That said, these are still very simplistic models that invite new modeling innovations. Progress beyond the useful benchmark of Nash equilibrium has been so rapid in recent years, however, that one can 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] 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.

[2] 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.

[3] Bhatt, 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.

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

[4] Bhatt, Meghana A., Terry Lohrenz, Colin F. Camerer, and P. Read Montague. 2012. "Distinct contributions of the amygdala and parahippocampal gyrus to suspicion in a repeated bargaining game." Proceedings of the National Academy of Sciences, 109(22): 8728-8733.

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

[6] 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.

[7] 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.

[8] Brocas, Isabelle, Juan D. Carrillo, Colin F. Camerer, and Stephanie W. Wang. Forthcoming. "Imperfect choice or imperfect attention? Understanding strategic thinking in private information games." Review of Economic Studies.

[9] 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.

[10] Bulow, Jeremy, and Paul Klemperer. 2002 ÒPrices and the WinnerÕs Curse.Ó RAND Journal of Economics, 33(1), 1–21.

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

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

[13] Camerer, Colin, Salvatore Nunnari, and Thomas R. Palfrey. 2012. "Quantal Response and Nonequilibrium Beliefs Explain Overbidding in Maximum-Value Auctions", Caltech Social Science Working Paper 1349.

[14] Camerer, Colin, and Alec Smith. 2012. "Cognitive Hierarchies and Emotions in Behavioral Game Theory." The Oxford Handbook of Thinking and Reasoning: 346.

[15] 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.

[16] Carrillo, Juan D., and Thomas R. Palfrey. 2011. "No Trade." Games and Economic Behavior, 71(1): 66-87.

[17] 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.

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

[19] 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.

[20] Coricelli, Giorgio, and Rosemarie Nagel. 2009. "Neural correlates of depth of strategic reasoning in medial prefrontal cortex." Proceedings of the National Academy of Sciences, 106(23): 9163-9168.

[21] 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.

[22] Costa-Gomes, Miguel A., and Makoto Shimoji. 2011. "A Comment on 'Can Relaxation of Beliefs Rationalize the Winner's Curse?: An Experimental Study." Working Paper, University of Aberdeen.

[23] Crawford, Vincent P., Miguel A. Costa-Gomes, and Nagore Iriberri. 2013. "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications." Journal of Economic Literature, 51(1): 5–62.

[24] 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.

[25] 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.

[26] 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

[27] Esponda, Ignacio and Emanuel Vespa. 2014. "Hypothetical Thinking and Information Extraction in the Laboratory." AEJ: Micro, forthcoming.

[28] Eyster, Erik, and Matthew Rabin. 2005. "Cursed equilibrium." Econometrica, 73(5): 1623-1672.

[29] 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.

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

[31] 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.

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

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

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

[35] 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.

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

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

[38] McKelvey, Richard D., and Thomas R. Palfrey. 1995. "Quantal Response Equilibria for Normal Form Games." Games and Economic Behavior 10(1): 6–38.

[39] 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.

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

[41] Niederle, Muriel. In press. "Intelligent Design: The Relationship of Economic Theory to Experiments: Treatment driven Experiments". In "Methods of Modern Experimental Economics", edited by Guillaume Frechette and Andrew Schotter, Oxford University Press (available at http://www.stanford.edu/∼niederle/Intelligent%20Design.August15.pdf).

[42] Polonio, Luca, Sibilla Di Guida, Giorgio Coricelli, Sibilla Di Guida, and Giorgio Coricelli. 2014. "Strategic sophistication and attention in games: an eye-tracking study." Unpublished.

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

[44] Roth, Alvin. 1995. "Bargaining Experiments" in The Handbook of Experimental Economics. J. Kagel and A. Roths, Eds., Princeton, NJ: Princeton University Press.

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

[46] Selten, Reinhard. 1967. "Die Strategiemethode zur Erforschung des Eingeschrankt Rationalen. Verhaltens im Rahmen eines Oligopolexperiments," in Beitrage zur Experimentellen Wirtschaftsforschung. H. Sauermann, Ed. , pp. 136-168, Tubingen: J. C. B. Mohr.

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

[48] 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.

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

[50] 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.

[51] 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.

[52] 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.

[53] Yoshida, Wako, Ben Seymour, Karl J. Friston, and Raymond J. Dolan. 2010. "Neural mechanisms of belief inference during cooperative games." The Journal of Neuroscience 30(32): 10744-10751.

Appendix - Comparison of Our Data and ILN Data

Figure 11 reports the unconstrained bids by signal for Phase I, for both datasets. All bids are shown on a log scale (since bids can range from 0 to $1,000,000$). There is one striking feature of ILN data which is not shared with our data: extreme overbidding. 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). Surprisingly, 27% of ILN's subjects bid 100 or more at least once. These surprisingly high bids were not observed in our data

Figure 11: Bids as a function of signals in our data (a) and ILN (2010) (b), Phase I. All bids.

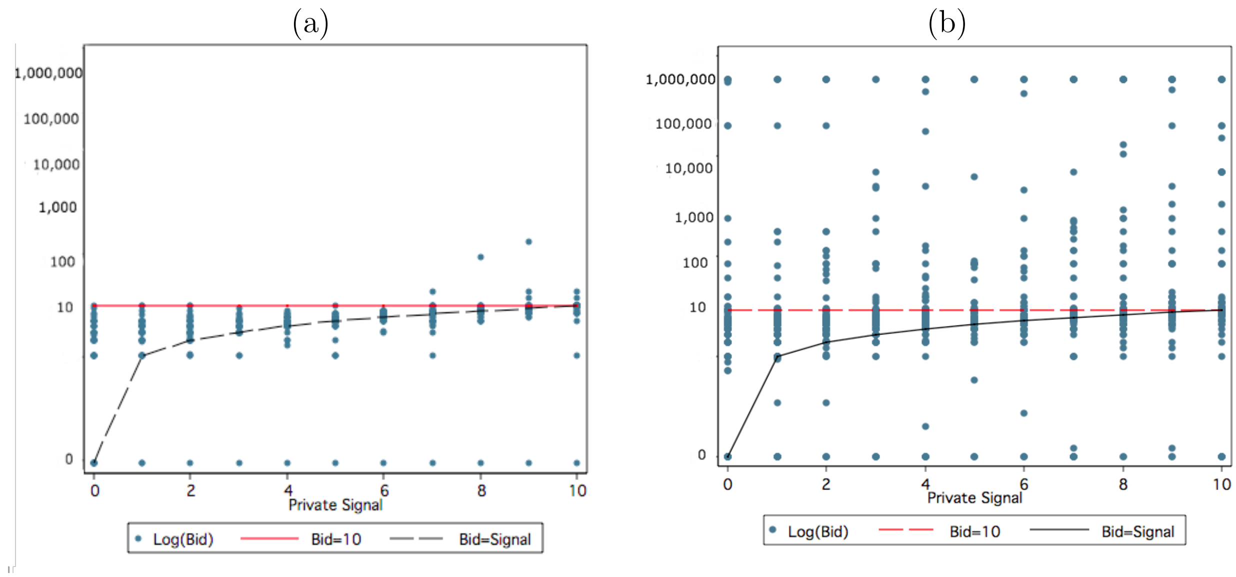

Table 3 and Figures 12-17 present the QRE, HQRE, CE-QRE, CE-HQRE, LK-QRE, and LK-HQRE predictions for the ILN data. In ILN, Phase II was conducted under two different protocols. In one, subjects were not reminded how they bid in Phase I. In the other protocol (ShowBid), subjects were reminded what bids they chose for each signal in Phase I. Our experiment uses only the ShowBid procedure to avoid concerns about subjects' memory. In the analysis of ILN data, we pool together data from these two protocols. The left panel of each figure shows predictions for Phase I obtained by estimating the parameters of the model as described in Section 3. The right panel of each figure, instead, shows the out-of-sample predictions for Phase II obtained using the parameters estimated with the data from Phase

I.

Figure 12: QRE Predictions, Baseline & ShowBid Treatment, ILN

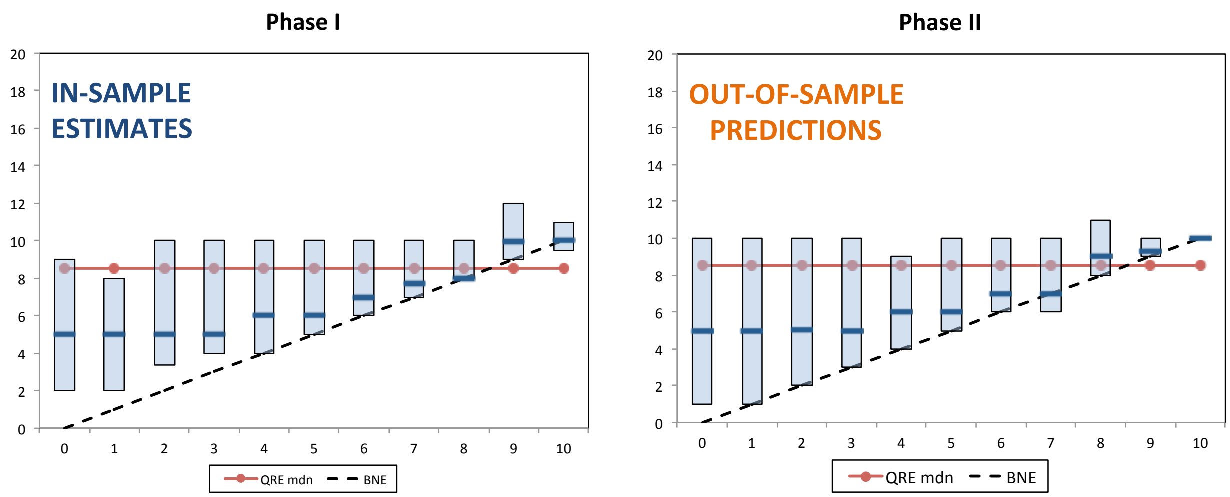

Figure 13: HQRE Predictions, Baseline & ShowBid Treatment, ILN

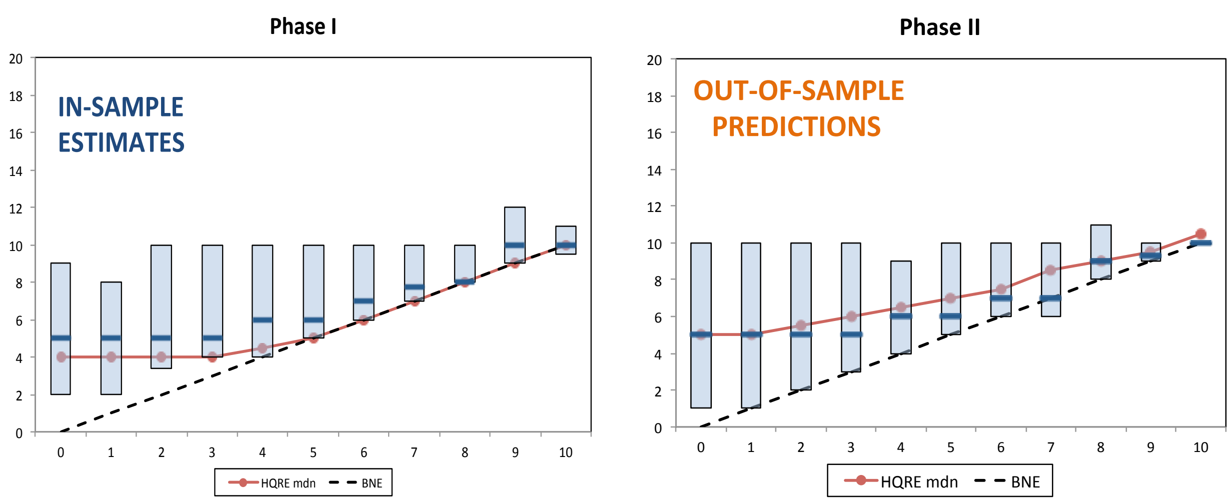

Table 3: Estimated Parameters

	ILN	**CNP**
# Individuals	108	46
# Observations	1188	506
QRE		
$\widehat{\lambda}$	0.00	3.59
-Log-likelihood (P1)	4223.8	1447.8
HQRE		
$\widehat{\lambda}_i$ Quartiles	(0.12, 1.30, 4.51)	(3.40, 5.12, 8.70)
-Log-likelihood (P1)	3582.0	1317.6
CE-QRE		
$\widehat{\chi}$	[0,1]	0.1
$\widehat{\lambda}$	0.00	3.66
-Log-likelihood (P1)	4223.8	1440.4
CE-HQRE		
$\widehat{\chi}$	0.5	0.2
$\widehat{\lambda}_i$ Quartiles	(0.11, 1.53, 6.76)	(3.05, 6.55, 11.81)
-Log-likelihood (P1)	3549.6	1289.7
LK-QRE (4 Types)		
$\widehat{\tau}$	0.8	2.7
$\widehat{\lambda}$	15.74	12.93
Types Frequency	(45%, 36%, 15%, 4%)	(9%, 25%, 34%, 31%)
-Log-likelihood (P1)	3895.3	1346.3
LK-HQRE (4 Types)		
$\widehat{\tau}$	2.9	4
$\widehat{\lambda}_i$ Quartiles	(0.12, 1.88, 18.05)	(3.82, 11.06, 47.51)
Types Frequency	(8%, 24%, 35%, 33%)	(4%, 17%, 34%, 45%)
-Log-likelihood (P1)	3525.8	1280.9

Figure 14: CE-QRE Predictions, Baseline & ShowBid Treatment, ILN

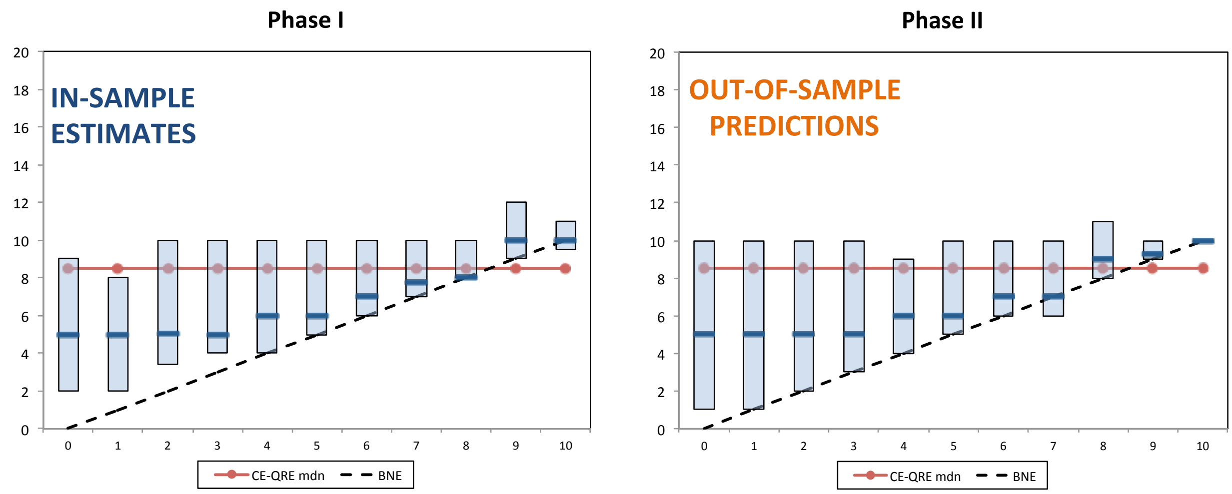

Figure 15: CE-HQRE Predictions, Baseline & ShowBid Treatment, ILN

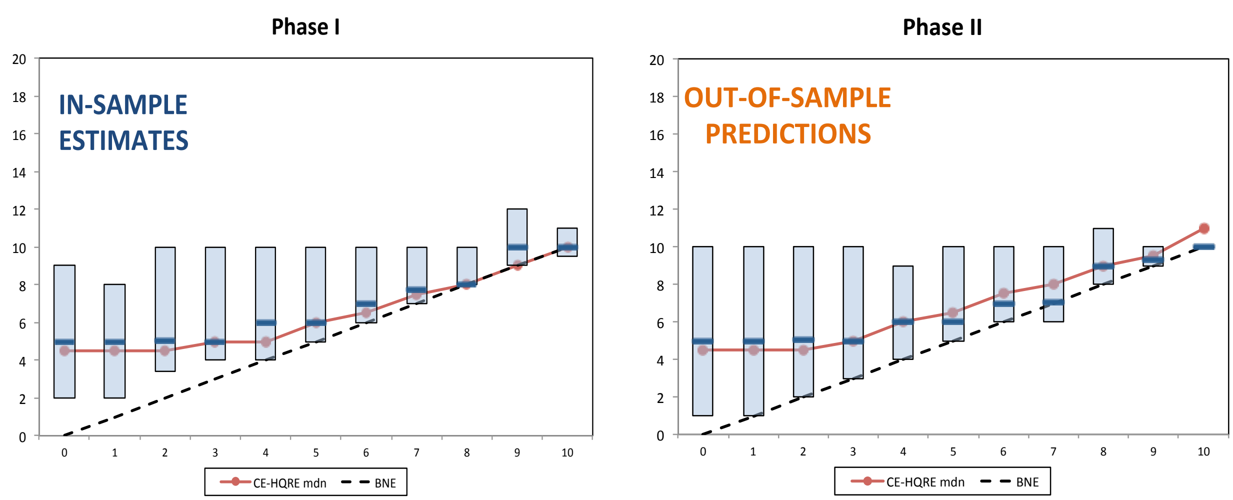

Figure 16: LK-QRE Predictions, Baseline & ShowBid Treatment, ILN

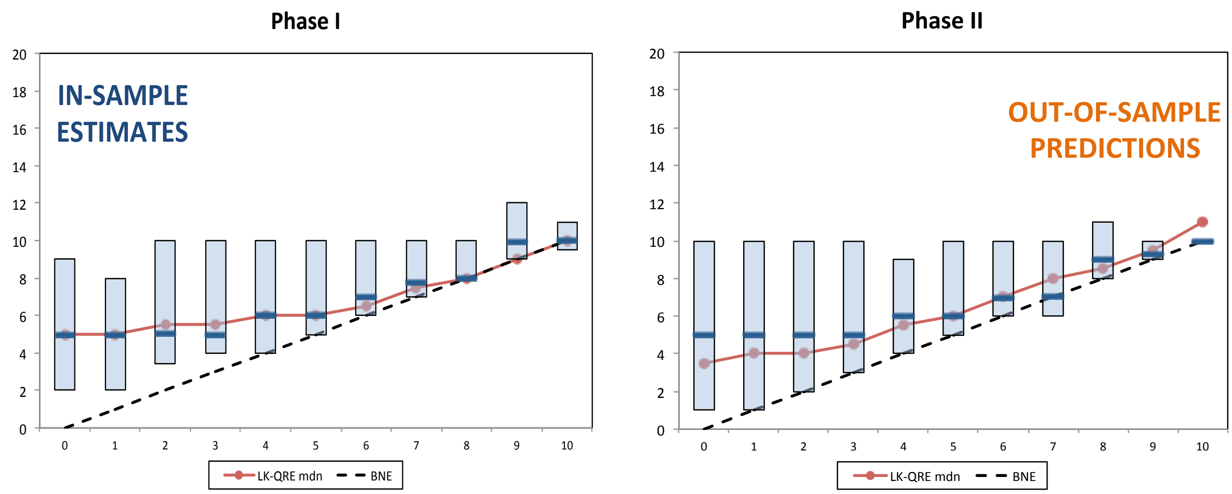

Figure 17: LK-HQRE Predictions, Baseline & ShowBid Treatment, ILN