



ELSEVIER

Contents lists available at ScienceDirect

Journal of Behavioral and Experimental Economics

journal homepage: www.elsevier.com/locate/jbee

Second-price common value auctions with uncertainty, private and public information: Experimental evidence[☆]



Isabelle Brocas^a, Juan D. Carrillo^{a,*}, Manuel Castro^b

^aDepartment of Economics, University of Southern California and CEPR, Los Angeles, CA 90089, United States

^bAdSupply, Inc., United States

ARTICLE INFO

Article history:

Received 1 August 2016

Revised 26 December 2016

Accepted 27 December 2016

Available online xxx

JEL classification:

C92

D44

D82

Keywords:

Laboratory experiments

Second-price common value auctions

Winner's curse

Level-k

Cursed equilibrium

ABSTRACT

We conduct a laboratory experiment of second-price sealed bid auctions of a common value good with two bidders. Bidders face three different types of information: common uncertainty (unknown information), private information (known by one bidder) and public information (known by both bidders), and auctions differ on the relative importance of these three types of information. We find that subjects barely differentiate between private and public information and deviate from the theoretical predictions with respect to all three types of information. There is under-reaction to both private and public information and systematic overbidding in all auctions above and beyond the standard winner's curse. The Cursed Equilibrium and Level-k models successfully account for some features of the data but others remain largely unexplained.

© 2017 Elsevier Inc. All rights reserved.

1. Introduction

Auctions are widespread but complex allocation mechanisms. They have received substantial attention in the theoretical and experimental literature in economics. Despite some departures from theoretical predictions, auctions are generally seen as reasonably efficient mechanisms for the allocation of items among agents. While many informational elements are present in an auction, the existing experimental literature generally focuses on one, the amount of private information. In this paper, we take a different route and study how subjects react to *three different types of information*: private information (information known by one bidder but not the other), public information (information known by both bidders) and common uncertainty (information known by no bidder). The main goal is to determine if subjects realize that optimal bidding depends not only on the informa-

tion they possess/ignore but also on whether this information is possessed/ignored by other bidders.

To this purpose, we adopt the experimental design first introduced by Brocas et al. (2015) in the context of a first-price sealed bid auction, add some small variants, and study behavior in a second-price sealed bid auction. Formally, we assume that the value of the good is the sum of N independent “components.” Each bidder observes only a subset of these components and knows which components are and are not observed by the other bidder. By varying the number of components observed by each bidder, we parsimoniously change the information structure of the auction. We consider five information structures. Three have only one type of information: only common uncertainty, only private information, and only public information. The remaining structures have two types of information: one has private information and common uncertainty and the other has private information and public information.

Brocas et al. (2015) show that more than half of the subjects in the first-price auction do not differentiate between types of information and that departures from equilibrium predictions occur with respect to all three types. As we will develop below, we also find that a majority of subjects in our second-price auction react similarly to private and public information. They also react differently from what theory predicts under all information struc-

[☆] We gratefully acknowledge the financial support of the National Science Foundation grant SES 1425062, the LUSK Center for Real Estate, the Microsoft Corporation and the Fundacao para a Ciencia e Tecnologia from the Portuguese Ministry of Science and the Consolider Program from the European Research Council.

* Corresponding author.

E-mail addresses: brocas@usc.edu (I. Brocas), juandc@usc.edu (J.D. Carrillo).

tures. Taken together, the two papers point to a general behavioral anomaly regarding information processing: subjects do not act optimally given their information because they do not realize that possessing information matters, nor that it matters whether or not their rival possesses that information. Overall, while the literature on adverse selection has extensively documented that informational asymmetries result in suboptimal choices (Samuelson and Bazerman, 1985; Charness and Levin, 2009; Carrillo and Palfrey, 2009; 2011), this and our companion paper are the first to explicitly show that inefficient choices occur because many subjects behave similarly when information is private or public.

A key distinguishing feature between the two studies is that the symmetric Nash Equilibrium (NE) bidding strategy in the second price auction can be decomposed into additively separable parts, each related to a type of information. Therefore, the individual effects of each type of information on the bidding strategy can be analyzed separately.¹ With respect to private information, subjects must bid their strategically estimated valuation (which is twice their signal as in the typical second-price common value auction à la Milgrom and Weber, 1982). With respect to common uncertainty and public information, subjects compete à la Bertrand and bid the expected value and the realized value, respectively. As a result, we can study the marginal effect of each change in the information structure on the bid. The separability property also lends itself to a structural estimation of behavioral models. To assess the behavior of subjects and compare it with the NE prediction, we analyze our experimental data from three different angles: we perform an aggregate descriptive analysis, we run a regression analysis to explain the bid as a function of the information known by the subject (public and private), and we perform a structural estimation of two behavioral theories, Level-k (Lk) and Cursed Equilibrium (CE).

The main conclusions of our analysis are the following. As mentioned above, the first and arguably most robust result is that the subjects' reaction to new information depends only marginally on its type. In other words and complementing Brocas et al. (2015), bidders treat private and public information much more similarly than they should. The second finding is that the observed behavior departs from NE regarding the three types of information and these departures depend on the information revealed. There is overbidding of common uncertainty which increases with total information. The reaction to private information is significantly smaller than predicted by theory and it increases with total information. Finally, the reaction to public information is constant and slightly smaller than NE. Significant departures occur even in auctions with no information and full information. The third conclusion relates to behavioral theories. Both the CE and Lk models can be seen as successful in that they parsimoniously explain some important features of the data, namely overbidding when the value of the private information is small and under-reaction to increases in that information. However, both models fail to capture the extra overbidding (above and beyond the winner's curse) and the substantial heterogeneity observed in our sample. It suggests that there is still room for improvement on existing theories.

Our analysis relates to two strands of the experimental literature: common value auctions and auctions with variable information. Second-price common value auctions have been extensively studied in the laboratory. In the typical setting (e.g., Kagel et al., 1995), the good is drawn from some distribution and bidders receive independent signals centered around the true realization. In our study, we model the value of the good as the sum of N independent signals, and each of them may or may not be

observed by bidders. This is formally closer to Albers and Harstad (1991); Avery and Kagel (1997) and Klemperer (1998).² As noted above, our paper is different in that we explicitly model components characterized by different types of information for which subjects bid independently, and we vary the relative importance of those components for comparative statics of bids and payoffs.

A few experimental articles study auctions with different amounts of information. Andreoni et al. (2007) study private value first- and second-price auctions in which bidders know their own valuation and the valuation of some other bidders. Naturally, the private value setting precludes any winner's curse problem. Mares and Shor (2008) analyze common value first- and second-price auctions with constant informational content but distributed among a varying number of bidders. The paper explores the trade-off competition vs. precision of estimates. Grosskopf et al. (2010) vary the number of bidders who receive a signal about the common value of the good in a first-price auction. They find that the winner's curse increases with private information. However, they do not consider how other types of information may affect the bidding strategy of the subjects.

The paper proceeds as follows. The theoretical framework is developed in Section 2 and the experimental setting and hypotheses are presented in Section 3. The aggregate analysis of the experimental data, including the regression analysis, is discussed in Section 4. Behavioral models are tested in Section 5 and conclusions are presented in Section 6.

2. Theoretical model

Consider a single common value good made of N components (with N even and greater than or equal to four). Each component $i \in \{1, \dots, N\}$ has a value x_i independently drawn from a continuous distribution with positive density $g(x_i)$ on $[\underline{x}, \bar{x}]$ and cumulative distribution $G(x_i)$. The total value of the good is the same for every individual and equal to the sum of the components, $V = \sum_{i=1}^N x_i$.

Two risk-neutral bidders, A and B indexed by j , bid for this good in a second-price sealed bid auction with no reserve price. Before placing their bids, A observes the first r components of the good, $\{x_1, \dots, x_r\}$, and B observes the last r components of the good, $\{x_{N-r+1}, \dots, x_N\}$, where $r \in \{1, \dots, N\}$. We also consider the case where bidders A and B observe none of the components of the good, which we denote by $r = 0$.

In this model, each bidder observes exactly r components and does not observe exactly $N - r$ components, and each bidder knows which components are and are not observed by the other bidder. We can define three types of information: *private information*, the components that only one bidder observes; *public information*, the components that both bidders observe; and *common uncertainty*, the components that no bidder observes. Notice that there is only common uncertainty when $r = 0$, only private information when $r = N/2$ and only public information when $r = N$. There is common uncertainty and private information when $r \in \{1, \dots, N/2 - 1\}$ and private information and public information when $r \in \{N/2 + 1, \dots, N - 1\}$. For the rest of the analysis, it is useful to introduce the following notation.

- $X_A^r = \sum_{i=1}^{\min\{r, N-r\}} x_i$: sum of A's private information when $r \in \{1, \dots, N - 1\}$
- $X_B^r = \sum_{i=\max\{N-r+1, r+1\}}^N x_i$: sum of B's private information when $r \in \{1, \dots, N - 1\}$

² In the first of these studies the value of the good is the sum of N signals and each of the N bidders observes one signal. In the last two studies, each of two bidders has one private signal. The value of the good is the sum of signals for one bidder and the sum of signals plus a private value component for the other bidder. When the private value component is zero, their model is equivalent to our treatment with only private information.

¹ By contrast, in a first-price auction, as information about certain components gets revealed, the bidding strategy changes for all the components of the good (and not only for the components affected by the change). It is therefore difficult to pinpoint the contribution of each component to the bidding strategy.

- $E[X_{\theta}^r] = \sum_{i=r+1}^{N-r} E[x_i]$: expected common uncertainty when $r \in \{0, \dots, N/2 - 1\}$
- $X_{\text{Pub}}^r = \sum_{i=N-r+1}^r x_i$: sum of public information when $r \in \{N/2 + 1, \dots, N\}$

Proposition 1 characterizes the symmetric Nash Equilibrium (NE) bidding strategies in this auction as a function of r .

Proposition 1. *The symmetric Nash Equilibrium bidding strategy of bidder j is:*

- $b_j^r = E[X_{\theta}]$ when $r = 0$,
- $b_j^r(X_j^r) = E[X_{\theta}] + 2X_j^r$ when $r \in \{1, \dots, N/2 - 1\}$,
- $b_j^r(X_j^r) = 2X_j^r$ when $r = N/2$,
- $b_j^r(X_j^r) = X_{\text{Pub}}^r + 2X_j^r$ when $r \in \{N/2 + 1, \dots, N - 1\}$,
- $b_j^r = X_{\text{Pub}}^r$ when $r = N$.

Proof. Standard. We restrict attention to differentiable monotonic bidding strategies. Assume that B bids in round r according to such a function and denote it by $b^r(X_B)$.

Let $r \in \{1, \dots, N/2 - 1\}$. The expected utility of A when he bids b_A^r is:

$$U_A^r = \Pr(b_A^r \geq b^r(X_B^r))(X_A^r + E[X_{\theta}^r] + E[X_B^r - b_B^r(X_B) \mid b_A^r \geq b^r(X_B^r)])$$

It can be rewritten as:

$$U_A^r = (X_A^r + E[X_{\theta}^r])F^r(b^{r-1}(b_A^r)) + \int_{X_B^r}^{b^{r-1}(b_A^r)} (X_B^r - b^r(X_B^r))f^r(X_B^r)dX_B^r$$

where $F^r(X_A^r) = \int_{x_1}^{x_2} \dots \int_{x_{r-1}}^{x_r} G(X_A^r - x_1 - \dots - x_{r-1})g(x_1) \dots g(x_{r-1})dx_1 \dots dx_{r-1}$.

Maximizing U_A with respect to b_A^r and imposing the symmetry condition $b_A^r = b^r$, yields the result. The proof for the other rounds follows the same lines. \square

The model is an extension of the “wallet game” to multiple types of information (see [Albers and Harstad, 1991](#); [Avery and Kagel, 1997](#), or [Klemperer, 1998](#)). In our model, the optimal bidding function can be split into two parts. The first part reflects common uncertainty when $r < N/2$ and public information when $r > N/2$, while the second part reflects private information for all $r \notin \{0, N\}$. For the first part, the risk-neutral agents compete à la Bertrand and end up bidding the expected value of the common uncertainty or the realized value of the public information. For the second part, agents bid ‘as if’ the private information of the opponent is at most equal to theirs, which happens to be true in the symmetric equilibrium. Overall, when $r = N/2$, the model is identical to the wallet game. When $r \notin \{0, N/2, N\}$, it can be seen as a wallet game with a third wallet whose content is known by either both bidders ($r > N/2$) or no bidder ($r < N/2$).

Notice that second-price common value auctions have multiple asymmetric equilibria. For example one subject bidding the maximum possible value, independently of the information, and the other subject bidding the lowest possible value, also independently of the information, is always an equilibrium of our game.³ In this paper, we focus on the symmetric equilibrium, although we will briefly discuss the problem of multiplicity when we analyze heterogeneity in the behavior of subjects.

3. The experiment

3.1. Design and procedures

We conducted 8 sessions with either 10 or 12 subjects per session for a total of 92 subjects. Subjects were undergraduate

³ Multiplicity of equilibria is discussed by [Milgrom \(1981\)](#) and [Avery and Kagel \(1997\)](#) in a general setting and by [Klemperer \(1998\)](#) in the context of the wallet game.

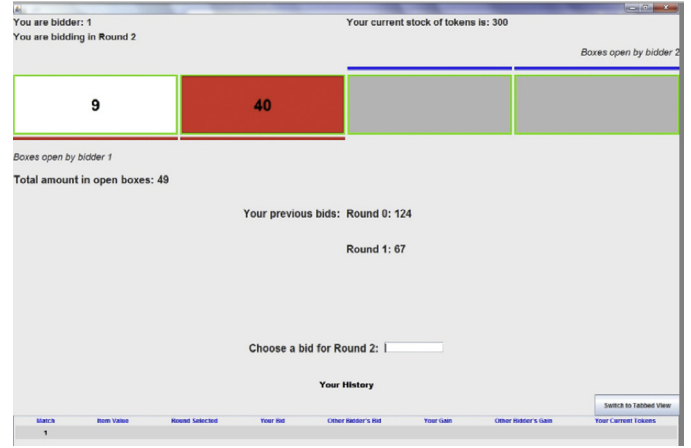


Fig. 1. User interface for the second-price auction.

students at the University of California, Los Angeles who were recruited by email solicitation, and all sessions were conducted at the California Social Science Experimental Laboratory (CASSEL). All interaction between subjects was computerized using an extension of the open source software package Multistage Games.⁴ No subject participated in more than one session.

In each session, subjects made decisions over 15 paid matches, with each match being divided into 5 rounds. At the beginning of a match, subjects were randomly matched into pairs and randomly assigned a role as bidder A or bidder B . Pairs and roles remained fixed for the 5 rounds of a match. At the end of the match, subjects were randomly rematched into new pairs and reassigned new roles.

The game closely followed the setting described in [Section 2](#). Subjects in a pair had to bid in a second-price sealed bid auction for a good made of $N = 4$ components. Each component $i \in \{1, \dots, 4\}$ contained x_i tokens drawn from a uniform distribution in $[0, 50]$ (to simplify computations, we restricted x_i to integer values). The total value of the good, V , was common to both bidders and equal to the sum of the four components, $V = \sum_{i=1}^4 x_i$. Visually, each component was represented by a box on the computer screen. The number of tokens inside each of the four boxes was drawn at the beginning of the match and did not change during the match. Subjects could see the four boxes but not their content.

The match was then divided into *five rounds*. Round 0 corresponded to $r = 0$ in the theory section, where bidders could not see the content of any of the boxes. Both subjects submitted a bid for *the entire good* of value $V = \sum_{i=1}^4 x_i$. Subjects could not see the bid of their rival, instead they moved to round 1. Round 1 corresponded to $r = 1$ in the theory section, where subjects A and B could observe x_1 and x_4 respectively and placed a new bid again for the entire good V . Again, subjects could not see the bid of their rival and moved directly to round 2, and so on until round 4 where both bidders could see the content of all 4 boxes. At the end of round 4, the value V of the item and the five bids of each subject were displayed on the computer screen. One of the rounds was randomly selected by the computer, and subjects were paid for their performance in that round.⁵ A sample screenshot of the user interface in round 2 is presented in [Fig. 1](#). It displays the subject’s role, the current round, the stock of tokens, the content

⁴ Documentation and instructions for downloading the software can be found at <http://multistage.ssel.caltech.edu>.

⁵ Notice that the order of play is always the same: from no information, to private information, to full information. Our experimental results may be affected by this specific order. However, we felt that this potential problem was offset by having a natural sequential presentation of information revelation.

of the open boxes, a reminder of which boxes have been opened by the rival and a reminder of the bid(s) of the subject in the previous round(s) of that match (but not the bid(s) of the rival).

Payoffs were computed according to the standard rules of a second-price auction without reserve price: the highest bidder won the item and paid the bid of the lowest bidder, while the lowest bidder got nothing and paid nothing. Rounds 0 and 4 (with no and full information) were included to facilitate comparisons across rounds and to better control for effects unrelated to private information (risk attitudes, joy of winning, etc.). Round 2 had only private information (the standard wallet game).

It is crucial to notice that in a second-price auction the optimal bid for each component x_i is *independent* of the other components (see [Proposition 1](#)). In other words, if x_1 and x_4 are private information for subjects A and B , these subjects should bid $2x_1$ and $2x_4$ respectively for components one and four, independently of whether components two and three are common uncertainty ($r = 1$), private information ($r = 2$) or public information ($r = 3$). This is key for our experiment as it implies that the optimal bid in round r can be decomposed into the optimal bid in round $r - 1$ plus the effect of the new information revealed in round r . It is a key distinction with the first-price auction ([Brocas et al., 2015](#)) where such decomposition is not possible.

All participants started the experiment with an endowment of 300 tokens to which the payoffs of each match were added or subtracted. To minimize bankruptcy situations, participants could not bid more than their current stock of tokens.⁶ We also constrained the bids to be between 0 and 200: the minimum and maximum possible values of the good before any information is revealed.⁷

At the beginning of each session, instructions were read by the experimenter standing on a stage in the front of the experiment room (a sample copy of instructions can be found in [Appendix A](#)). The experimenter fully explained the rules and how to operate the computer interface. After the instructions were finished, one practice match consisting of all five rounds was conducted, for which subjects received no payment. After the practice match, there was an interactive computerized comprehension quiz that all subjects had to answer correctly before proceeding to the paid rounds. Then, the 92 subjects participated in 15 paid matches each of them divided into 5 rounds for a total of 75 bids per subject. Opponents, roles and values in all four boxes were randomly reassigned at the beginning of each match and held constant between rounds of a match. In the end, subjects were paid, in cash, in private, their accumulated earnings, which was equal to their initial endowment plus the payoffs of all matches. The conversion rate was \$1.00 for 20 tokens, so each good was worth between 0 and \$10. Sessions averaged 75 minutes in length, and subjects earnings averaged \$18 plus a \$5 show-up fee.

3.2. Hypotheses

In this section we briefly discuss our three hypotheses regarding the behavior of subjects in the experiment.

Hypothesis 1. Subjects sharply differentiate between public and private information.

⁶ Subjects in latter matches could have less tokens than necessary to play the Nash equilibrium and therefore be exogenously constrained. Because we gave enough initial capital, this happened in our experiment in less than 1% of the observations, having no significant effect in our analysis (for a study of selection bias in repeated auctions, see [Casari et al., 2007](#)).

⁷ The obvious reason for such constraint is to avoid extreme choices with large payoff consequences. It could potentially bias our results towards underbidding when the value of the item was close to 200. However, it did not have a significant effect in the analysis.

Perhaps the most novel feature of the experiment is the manipulation of the amount of public and private information, without affecting all the other elements of the auction (the value of the good, the opponent, the complexity of the auction, etc.). To this purpose, opening one box at a time seems a particularly suitable methodology. [Hypothesis 1](#) states that subjects will behave very differently if the new information is known to be privately or publicly revealed, independently of whether they play in each case according to Nash theory or not. To our knowledge this paper and [Brocas et al. \(2015\)](#) are the first experiments designed to address the effect on bidding behavior of “knowing what the other knows”, holding everything else constant.

Hypothesis 2. Subjects deviate from Nash equilibrium predictions with respect to all three types of information. However, deviations will be more substantial with respect to private information than with respect to public information and common uncertainty.

Note that [Hypothesis 1](#) does not claim that subjects will closely follow the theoretical predictions. Indeed, we know from previous experiments that subjects fall prey of the winner's curse in common value games with private information ([Kagel and Levin, 2008](#) and others) and over-price in Bertrand competition games ([Abbink and Brandts, 2008](#) and others). [Hypothesis 2](#) states that, in accordance to previous research, we expect deviations with respect to all three types of information. However, our design will allow us to determine whether deviations from equilibrium bids are more pronounced with respect to one type of information or another, and also which shape they take (overbidding or underbidding and overreaction or under-reaction to information).

Hypothesis 3. Cursed equilibrium and level-k theories can explain the behavioral departures of our subjects.

Recent behavioral theories (cursed equilibrium, level-k and other combinations) have received empirical support in private information experiments in general ([Carrillo and Palfrey, 2009](#) and [Rogers et al., 2009](#)) and in auctions in particular ([Eyster and Rabin, 2005](#) and [Crawford and Iriberri, 2007](#)). [Hypothesis 3](#) argues that a parametric estimation of these two behavioral theories should fit the data reasonably well. Notice that this hypothesis is indirectly related to [Hypothesis 2](#). Since cursed equilibrium predicts no deviation with respect to public information and common uncertainty, a good fit is only possible if the major driving departure from equilibrium behavior is due to private information.

4. Aggregate analysis

Our first analysis consists in comparing the aggregate results in our sample with the NE predictions. To do this, we construct a new sample containing the predicted bids if everyone played the symmetric NE strategy. In the first part of this section we present a simple unconditional aggregate analysis of bids and payoffs in each round. In the second part, we study how the aggregate bids change over rounds in order to understand how subjects perceive the different types of information. In the third part, we use an OLS regression to study the empirical bidding as a function of the subject's information, both private and public.

4.1. Aggregate bids and payoffs

The first cut at the data consists in an aggregate analysis per round in order to compare actual behavior with the NE predictions derived in [Proposition 1](#). [Fig. 2](#) shows the difference between actual bids and NE predictions in each round $r \in \{0, 1, 2, 3, 4\}$. For each observation, we compute the NE bid and subtract it from the corresponding observation. The line in the middle is the median of this statistic, whereas the top and bottom lines are the 75th and




Fig. 2. Deviations from Nash Equilibrium.

Table 1
Average bids.

Round	0	1	2	3	4
Mean Data	118.1*** (1.10)	116.2*** (1.08)	113.1*** (1.12)	110.8*** (1.15)	110.4*** (1.24)
Mean NE	100 (0.00)	101.05 (0.80)	99.76 (1.13)	99.76 (0.97)	99.76 (0.79)

*, **, ***: Significantly different from Mean NE at 90%, 95% and 98% confidence level (*t*-test) (Standard errors in parenthesis)

25th percentiles. The notches are the 95% confidence interval for the median. In rounds 1, 2 and 3 (auctions with different degrees of private information), we observe a large dispersion in the data, consistent overbidding with a median overbid of 10 tokens (that is, 10% above the median value), and a roughly symmetric pattern around the median. In round 0 (auction of an item with unknown value) there is also overbidding and dispersion, though quite asymmetric. Few subjects underbid, suggesting that risk-aversion is unlikely to play a major role in the subject's strategy. In round 4 (auction of an item with known and identical value for both bidders), there is also some aggregate overbidding, fewer cases of underbidding than in round 0, and more than half of the bids at or close to NE. It suggests that joy of winning and other psychological factors may account for some but not a large part of the subject's behavior.

Table 1 displays the average bids per round (Mean Data) and the equilibrium predictions (Mean NE). As in any standard second-price auction, the unconditional NE average bid is equal to the unconditional expected value of the good (which is different from 100 only because of natural variation in sample realizations). A comparison with Fig. 2 reveals that the average overbidding is higher than the median overbidding. This indicates the presence of outliers consisting of extremely high bids. As we will see later, some subjects bid the maximum amount of 200 tokens independently of their value and information. Moreover, the average bid decreases across rounds, so that more total information makes bidders (weakly) less aggressive.

Table 2 shows the average empirical and predicted gains. If all subjects follow the equilibrium strategy, they only extract rents from private information, since they compete à la Bertrand for the known and unknown components. Therefore, the predicted gains are hump-shaped across rounds: zero in rounds 0 and 4, and maximum in round 2. By contrast, the observed overbidding implies that gains in our sample are substantially smaller than NE in the rounds with private information (1, 2 and 3). In rounds 0

Table 2
Average gains.

Round	0	1	2	3	4
Mean Data	1.27* (0.83)	2.44*** (0.76)	4.53*** (0.67)	4.65*** (0.60)	3.49*** (0.61)
Mean NE	-0.12 (0.56)	7.87 (0.50)	12.17 (0.47)	8.52 (0.32)	0.00 (0.00)

*, **, ***: Significantly different from Mean NE at 90%, 95% and 98% confidence level (*t*-test) (Standard errors in parenthesis)

and 4, bidders have positive gains on average. This occurs because of a well-known problem in a second-price auction: if a few subjects bid at (or close to) zero, their rivals greatly benefit at the expense of the seller.

These preliminary findings are summarized in the following result.

Result 1. There is overbidding relative to NE and a large dispersion of bids in all rounds. Gains are smaller than NE in rounds with private information and larger otherwise.

4.2. Aggregate change in bids

The previous analysis is instructive but incomplete as it does not distinguish by amount or type of information. This section studies, still at the aggregate level, how subjects react to different types of information. A first natural question, raised in Hypothesis 1, is whether subjects understand the differences in the types of information and act accordingly. In rounds 1 and 2, the last box revealed consists of private information, whereas in rounds 3 and 4, the last box revealed consists of public information. Therefore, NE predicts that for the same value in the last box, the change in bids in rounds 1 and 2 should be different than the change in bids in rounds 3 and 4. In this section we test this prediction. Moreover, since bids in early rounds depart from equilibrium predictions, studying changes in bids across rounds also addresses the following question: are deviations in later rounds due to imperfect adjustments or to carrying over some initial miscalculations?

We analyze the change in bids between two consecutive rounds as a function of two variables: the value of the new box revealed and the sum of the open boxes. Formally we organize the two variables somewhat arbitrarily into 'high' values and 'low' values. High values (*H*) correspond to cases in which the sum of the values in the boxes already open (first variable) and the value in the new box (second variable) are above their expected amounts.⁸ Similarly, low values (*L*) correspond to cases in which the sum of the values in the boxes already open (first variable) and the value in the new box (second variable) are below their expected amounts. We construct four groups '*H* to *H*', '*H* to *L*', '*L* to *H*' and '*L* to *L*' where the first letter refers to the value of the sum of open boxes and the second to the value of the new box. For each of these four groups we calculate the average change in bids between round *r* and round *r* + 1 with *r* ∈ {0, 1, 2, 3} (bid in *r* + 1 minus bid in *r*), which we called '*r* to *r* + 1', and present the results in the columns of Table 3. The rows are divided into the four groups previously described. For each group we present the NE predictions and the data. The last two columns contain the *p*-value of a normal (*N*) and non-parametric (NP) test for the difference in means across rounds.

⁸ We use this crude distinction between 'high' and 'low' to have few categories with enough observations each. Similar results are obtained with three categories. With more categories the analysis lacks statistical power. Obviously, there are other ways to look at the data. For example, looking separately at cases where new (private or public) information requires an upward vs. a downward revision of the bid.

Table 3
Average change in bids over rounds (bid in $r + 1$ - bid in r).

		0 to 1 [§]	1 to 2	2 to 3	3 to 4	p-val \mathcal{N}^\dagger	p-val NP ^{††}
L to L	NE	-25.61	-25.44	-3.68	-6.96	0.00	0.00
	Data	-12.43	-15.49	-15.37	-12.69	0.30	0.49
H to L	NE		-24.65	-20.17	-19.78	0.00	0.00
	Data		-12.05	-10.67	-10.93	0.79	0.78
L to H	NE	25.85	26.53	22.53	18.94	0.00	0.00
	Data	7.70	10.48	7.87	9.82	0.40	0.01
H to H	NE		25.30	4.94	5.54	0.00	0.00
	Data		7.94	11.66	11.04	0.20	0.04

§: Round 0 has no information so only the second category is relevant. †: ANOVA test. ††: Kruskal–Wallis test.

According to the NE described in Proposition 1, bids of subject A should be:

Round 0	Round 1	Round 2	Round 3	Round 4
$4E[x]$	$2E[x] + 2x_1$	$2x_1 + 2x_2$	$2x_1 + x_2 + x_3$	$x_1 + x_2 + x_3 + x_4$

This implies that the change in bids should be:

$$\begin{array}{cccc} 0 \text{ to } 1 & 1 \text{ to } 2 & 2 \text{ to } 3 & 3 \text{ to } 4 \\ 2x_1 - 2E[x] & 2x_2 - 2E[x] & x_3 - x_2 & x_4 - x_1 \end{array} \quad (1)$$

In the first two cases, the change in bids depends on private information and uncertainty whereas in the last two cases the change in bids depends entirely on public information (the behavior of B is determined analogously).

As we can see from the NE rows of Table 3, the change in bids predicted by theory depend (obviously) on whether the new and open boxes are above or below average. However, they also strongly depend on whether the new information is public or private. This can be best grasped by looking at the first and last groups.

In the 'L to L' group, the theoretical changes in the private information columns ('0 to 1' and '1 to 2') are $2E(x_1 | x_1 < E[x]) - 2E[x] = -25$ and $2E(x_2 | x_2 < E[x]) - 2E[x] = -25$ whereas the theoretical changes in the public information columns ('2 to 3' and '3 to 4') are $E(x_3 | x_3 < E[x]) - E(x_2 | x_2 < E[x]) = 0$ and $E(x_4 | x_4 < E[x]) - E(x_1 | x_1 < E[x]) = 0$ (the difference between these values and those reported under NE are due to the natural variation in sample realizations). In the 'H to H' group, the results are identical with opposite signs. Overall, in these two extreme cases, NE bid changes are large and significant under revelation of private information and insignificant under revelation of public information.

In the data, the change in bids between the private and public information columns is not statistically significant in 'L to L'. It is significant under the non-parametric test in 'H to H' but small and with the opposite sign than predicted by theory. So, in both groups, the difference in the reaction by our subjects to private vs. public information is not as sharp as the theory would predict. This is also the case in 'H to L' and 'L to H', although the theoretical differences are not as big, so the contrast is not as striking. All in all, subjects in our experiment act as if they only considered the amount in each box and not the type of information the box is revealing to themselves and to their opponents.⁹ This finding is summarized in the next result.

Result 2. Hypothesis 1 is not supported by the data: subjects barely distinguish between public and private information.

⁹ As developed in footnote 8, this is only one possible way of looking at the data across rounds. An alternative would be to run the regressions $y_1 = \alpha_0 + \alpha_1 x_1 + \epsilon_1$, $y_2 = \alpha_0 + \alpha_2 x_2 + \epsilon_2$, $y_3 = \alpha_0 + \alpha_2 x_2 + \alpha_3 x_3 + \epsilon_3$, $y_4 = \alpha_0 + \alpha_1 x_1 + \alpha_4 x_4 + \epsilon_4$ where y_r is the empirical change in bid between round $r - 1$ and round r . We could then compare the α -coefficients to those predicted by (1). This would be analogue to the within-round regression conducted in Section 4.3. Each method has its pros and cons.

4.3. Regression analysis of bidding strategies

In Section 4.1, we showed consistent overbidding in every round and in Section 4.2 we emphasized departures from theoretical predictions in the reaction to information *between* rounds. We now study how bids react to information *within* rounds. Indeed, the insufficient distinction between private and public information highlighted in Result 2 may be caused by deviations from Nash with respect to some or all types of information. To study this question, we run the following OLS regressions for rounds 0 to 4:

$$\begin{aligned} b_o^0 &= \alpha^0 + \epsilon_o^0 \\ b_o^1 &= \alpha^1 + \beta^1 \text{Priv}_o^1 + \epsilon_o^1 \\ b_o^2 &= \alpha^2 + \beta^2 \text{Priv}_o^2 + \epsilon_o^2 \\ b_o^3 &= \alpha^3 + \beta^3 \text{Priv}_o^3 + \gamma^3 \text{Pub}_o^3 + \epsilon_o^3 \\ b_o^4 &= \alpha^4 + \gamma^4 \text{Pub}_o^4 + \epsilon_o^4 \end{aligned}$$

In these equations, Priv_o^r is the variable of private information (tokens observed by only one bidder) in observation o of round r . This variable appears only in the OLS regression of the rounds where there is private information (1, 2 and 3). Pub_o^r is the variable of public information (tokens observed by both bidders) in observation o of round r . This variable appears only in the OLS regression of the rounds where there is public information (3 and 4). Finally, the constant term captures the common uncertainty (expected tokens in the components observed by no bidder). The predictions of Nash theory for the constant term (α^r parameter) are $\alpha^0 = 100$, $\alpha^1 = 50$, and $\alpha^2 = \alpha^3 = \alpha^4 = 0$, which correspond to the expected number of unobserved tokens in each round. The predictions of Nash theory for the coefficient of private information (β^r parameter) are $\beta^1 = \beta^2 = \beta^3 = 2$. Indeed, in the symmetric equilibrium of a second-price auction, agents bid as if the private information of the opponent is at most equal to theirs, hence the coefficient of the rival's private information is the same as the coefficient of the bidder's own private information. Finally, the predictions of Nash theory for the coefficient of public information (γ^r parameter) are $\gamma^3 = \gamma^4 = 1$, since there is Bertrand competition for the observed tokens.

The goal of the regression is to compare the α^r , β^r and γ^r coefficients estimated in our sample with the NE predictions using a t-test. Notice that, in contrast with some other auctions (and, in particular, with the first-price auction in Brocas et al., 2015) theoretical bids are linear in the three types of information, so no approximation is necessary. The coefficients have simple predictions and sharp interpretations. Furthermore, behavior across rounds can be easily compared because of the separability of bids with respect to components. Indeed, if a component keeps the same type of information in two different rounds, the bid for that component should not change.

The results of the five OLS regressions are displayed in the columns of Table 4. For each round, we include the relevant variable(s) (Priv., Pub., constant) and present both the theoretical

Table 4
Nash Equilibrium and OLS regression per round.

	Round 0		Round 1		Round 2		Round 3		Round 4	
	NE	Data	NE	Data	NE	Data	NE	Data	NE	Data
Priv.	—	—	2	0.75*** (0.002)	2	0.84*** (0.002)	2	0.95*** (0.002)	—	—
Pub.	—	—	—	—	—	—	1	0.86*** (0.002)	1	0.89*** (0.00)
constant	100	118.14*** (0.07)	50	96.99*** (0.10)	0	71.45*** (0.11)	0	44.44*** (0.13)	0	22.03*** (0.12)
F-test	274.15†††		266.43†††		377.47†††		130.29†††		59.68†††	
Adj. R ²	—		0.01		0.12		0.23		0.27	

Standard errors clustered by subject in parenthesis *, **, ***: Significantly different from NE at 90%, 95% and 98% confidence level (t-test) †, ††, †††: Significantly different from NE at 90%, 95% and 98% confidence level (F-test)

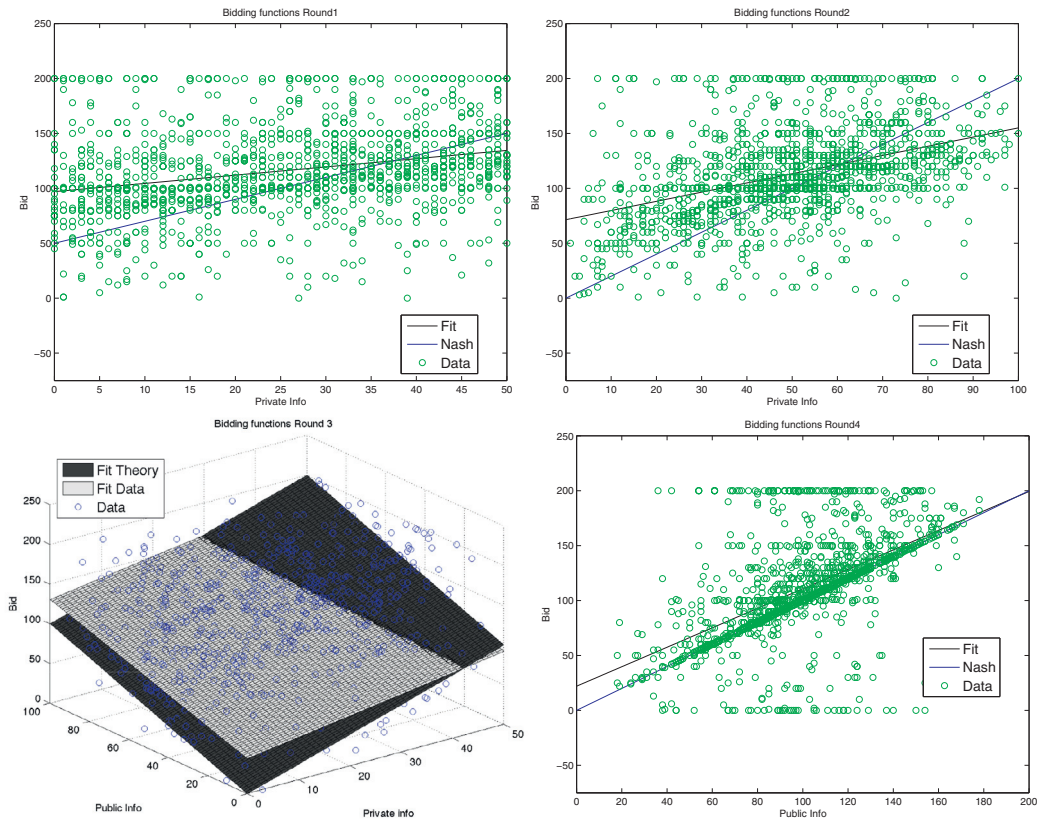


Fig. 3. NE and OLS regression: graphical display.

predictions (NE) and the estimated coefficients (Data). Fig. 3 presents all the bids in rounds 1–4 as a function of the information possessed by the bidder, whether it is only private (rounds 1 and 2, upper left and upper right graphs), only public (round 4, lower right graph) or both private and public (round 3, lower left three-dimensional graph). We also include for each graph the NE function and the OLS fit given by the estimates of Table 4.

We can see from Table 4 that all coefficients are individually and jointly significantly different from the Nash predictions. There is a large dispersion in the observations (as evidenced by Fig. 3) causing the adjusted R² in Table 4 to be small. Dispersion decreases with the amount of total information revealed, but remains surprisingly large even under full information (round 4). This may be partly due to some subjects playing the asymmetric equilibrium. Indeed, we can notice in the lower right graph of Fig. 3 a Z-shaped bidding function, with a significant fraction of bids equal to the lowest value independently of the information

($b_i(X_{Pub}) = 0$), others equal to the realized public information ($b_i(X_{Pub}) = X_{Pub}$), and finally some others equal to the highest value also independently of the information ($b_i(X_{Pub}) = 200$).

The constant parameter is always higher than NE, suggesting that overbidding is partly due to an over-estimation of the value of common uncertainty. Notice that if subjects neglected the fact that some of the closed boxes are observed by the rival (as suggested by Result 2) and the bidding of common uncertainty was constant over rounds, then the constant term in rounds 1, 2, 3 and 4 would be respectively 3/4, 1/2, 1/4 and 0 times the bid in round 0 (that is, 88.6, 59.1, 29.5 and 0, respectively). The estimated α^r coefficients in those rounds are significantly above those predictions but the difference is smaller than compared to the NE prediction. Overall, there is overbidding regarding common uncertainty, which increases with the total information revealed.

Subjects react much less to private information than what NE predicts, although it slightly increases with the amount of total

information revealed. The coefficients are all below 1 rather than the prediction of 2. Subjects also react less to public information than predicted by theory, but in this case the absolute difference between the estimated coefficients and the theoretical predictions is small. In a nutshell, bidding as a function of private and public information has a higher intercept and lower slope than the NE prediction. This is quite transparent in Fig. 3. Notice that the coefficients of private and public information are similar to each other which supports our previous finding that many subjects react similarly to private and public information. Last, notice that round 2 is equivalent to the auction treatment in Avery and Kagel (1997), but the results are somewhat different: subjects in our sample bid more aggressively (higher intercept) and react less to information (lower slope) than in theirs (in the next section we will compare the estimates of the level-k model across the two datasets). The differences could be due to anchoring effects, that is, to bids in rounds 0 and 1 having an effect on bids in round 2.¹⁰ The conclusions of this section are summarized in the following result.

Result 3. Hypothesis 2 is supported by the data: subjects depart from NE with respect to all three types of information. There is overbidding of common uncertainty, under-reaction to public information and strong under-reaction to private information. Deviations are least significant for common uncertainty.

5. Behavioral models

As shown above and perhaps not surprisingly, subjects in our experiment follow different strategies than NE. Several leading behavioral models have been proposed to explain departures in settings with private information (see e.g. Crawford and Iriberri (2007); Carrillo and Palfrey (2009) or Rogers et al. (2009) for estimations of several models within a given game). In this section, we explore two of them: Cursed Equilibrium (CE) and level-k (Lk). We focus on these theories because they are based on the idea that subjects fail to fully realize the effect of others' information on their behavior. Given the results of Section 4, these models are then good candidates to describe the strategies of our subjects. Naturally, it would be interesting to explore some other behavioral models, or combination of models, as well. Notice that the possibility of splitting the bidding function into independent components also extends to the behavioral models, which facilitates their structural estimation.

5.1. Theory 1: Cursed Equilibrium (CE)

In the CE model, each bidder systematically underestimates the correlation between the opponent's bid and the opponent's private information. In a χ -cursed equilibrium, each bidder plays a best response to beliefs that with probability χ her rival's action does not depend on her type, while with probability $1 - \chi$ her action does depend on her type, where $\chi \in [0, 1]$. The model is equivalent to NE when $\chi = 0$. Subjects are said to be "fully cursed" when $\chi = 1$. In our setting and following Eyster and Rabin (2005), the expected utility of a cursed bidder can be computed analytically. Let's denote by $b_j^{x,r}$ the bid of a χ -cursed subject j in round r . The next proposition characterizes the bid for subject A (the analysis for B is analogous).

Proposition 2. The difference between the CE bid and NE bid of bidder A is:

$$b_A^{x,r} - b_A^r = \chi (E[X_B] - X_A) \quad \text{if } r \in \{1, \dots, N-1\}$$

$$b_A^{x,r} - b_A^r = 0 \quad \text{if } r \in \{0, N\}$$

¹⁰ It would be also interesting to conduct an experiment explicitly designed to study this possibility.

This implies that cursed bidders overbid when their private information is below average and underbid when their private information is above average.

Proof. It follows Eyster and Rabin (2005). Let $r \in \{1, \dots, N/2 - 1\}$, assume B bids according to an increasing bidding function $b^{x,r}(X_B^r)$, and denote A 's bid by $b_A^{x,r}$. The expected utility of A can be written as:

$$U_A^{x,r} = \Pr(b_A^{x,r} \geq b^{x,r}) \left[\chi (X_A^r + E[X_B^r]) + E[X_B^r] \right. \\ \left. + (1 - \chi) (X_A^r + E[X_B^r] + E[X_B^r | b_A^{x,r} \geq b^{x,r}]) \right. \\ \left. - E[b_B^{x,r} | b_A^{x,r} \geq b^{x,r}] \right]$$

The CE bid can then be determined using the same procedure as we did for the NE bid, and we get:

$$b_A^{x,r} = E[X_B^r] + \chi E[X_B^r] + (2 - \chi) X_A^r$$

When $r \in \{N/2 + 1, \dots, N - 1\}$, we just need to replace $E[X_B^r]$ by X_{Pub}^r . When $r = 0$ or $r = N$, the CE bid coincides with the NE bid because there is no private information, and CE does not predict any deviations relative to the two other types of information. \square

Cursed bidders do not fully realize that the bid of the opponent is linked to his signal regarding the value of the item. Therefore, they react less to private information than NE bidders: the CE bidding function has a higher intercept and smaller slope than the NE bidding function. By contrast, since CE is a theory purely based on imperfect account of other's private information, it predicts no deviations from theory with respect to public information and common uncertainty (so, in particular, it coincides with NE in rounds 0 and N). Also, notice that because the CE bidding function is equal to NE when $X_A = E[X_B]$, the unconditional CE average bid is equal to the unconditional NE average bid. Finally, deviations from NE are stronger for more extreme (upward or downward) realizations of the signals, which are more likely to occur when the support of the distribution of private information is larger (that is, in round 2 of our experiment).

5.2. Theory 2: level-k (Lk)

The level-k model relaxes the assumption of accurate and homogeneous beliefs. It assumes the existence of different levels of strategic thinking in the population, and that each level best responds to lower levels. The model is built around an anchoring, non-strategic type: level-0. In this paper, we follow the approach developed by Nagel (1995) and Stahl and Wilson (1995) and applied to the auction setting by Crawford and Iriberri (2007), where each type believes everyone else's type corresponds to the level immediately below. Formally, level-1 best responds to level-0, level-2 best responds to level-1, and so on.¹¹

The definition of level-0 is crucial since it anchors the beliefs and actions of all other types. Following a common approach in level-k models, we assume that level-0 chooses a bid uniformly random over all possible bids.¹² As in Crawford and Iriberri (2007), we do not consider level-3 and above but we consider a last level-

¹¹ Instead, in the Poisson Cognitive Hierarchy theory (Camerer et al., 2004) each level believes that the population is distributed among all levels below.

¹² Crawford and Iriberri (2007) have an alternative definition of level-0 which they call 'Truthful.' However, as they discuss in the paper, in regards to private information the truthful level-0 corresponds to the level-1 of the hierarchy anchored on a random level-0 bidder. Thus, we adopted the more common definition of 'random level-0.' Moreover, note that a truthful level-0 treats public information and common uncertainty just like a NE player. Therefore, he would seem excessively sophisticated for a person who is usually considered "simplistic".

called *Equilibrium (Eq)*– consisting of NE bidders.¹³ Denote by \bar{X}_j^r and \underline{X}_j^r the maximum and minimum realization of subject j 's random variable X_j^r . Bidder A 's bidding function under the level- k theory is characterized below (again, the analysis for bidder B is analogous).

Proposition 3. *The bid of bidder A in level- k when $r \in \{1, \dots, N/2 - 1\}$ is:*

L_0	L_1	L_2	Eq
$U[0, 200]$	$X_A^r + E[X_\emptyset^r] + E[X_B^r]$	$\bar{X}_B^r + E[X_\emptyset^r] + E[X_A^r]$ if $X_A^r \geq E[X_A^r]$ $\underline{X}_B^r + E[X_\emptyset^r] + E[X_A^r]$ if $X_A^r < E[X_A^r]$	$E[X_\emptyset^r] + 2X_A^r$

When $r \in \{N/2 + 1, \dots, N - 1\}$, we replace $E[X_\emptyset^r]$ by X_{Pub}^r . When $r = N/2$, $E[X_\emptyset^r]$ is removed. When $r \in \{0, N\}$ all levels except L_0 bid like NE.

Proof. Let $r \in \{1, \dots, N/2 - 1\}$ (if $r \in \{N/2 + 1, \dots, N - 1\}$ then we replace $E[X_\emptyset^r]$ by X_{Pub}^r and if $r = N/2$ then $E[X_\emptyset^r]$ is removed).

- Level-1. L_0 bids uniformly random on the interval $[0, 200]$. Let $H(b) = b/200$. We can write the utility of a L_1 bidder A as:

$$U_A^r = (X_A^r + E[X_\emptyset^r] + E[X_B^r] - E[b_B^r | b_B^r \leq b_A^r])H(b_A^r)$$

$$\Leftrightarrow U_A^r = (X_A^r + E[X_\emptyset^r] + E[X_B^r])H(b_A^r) - \int_0^{b_A^r} b_B^r h(b_B^r) db_B^r$$

Taking the first-order condition with respect to b_A^r yields:

$$b_A^r(X_A^r) = X_A^r + E[X_\emptyset^r] + E[X_B^r]$$

- Level-2. Analogously, we can write the utility of a L_2 bidder A as:

$$U_A^r = (X_A^r + E[X_\emptyset^r] + E[X_B^r - b_B^r | b_A^r \geq b_B^r]) \Pr(b_A^r \geq b_B^r)$$

Substituting b_B^r by the bidding equation of a L_1 bidder B , we get:

$$U_A^r = (X_A^r - E[X_A^r])F^r(b_A^r - E[X_\emptyset^r] - E[X_A^r])$$

where $F^r(X_A^r) = \int_{\underline{x}}^{\bar{x}} \dots \int_{\underline{x}}^{\bar{x}} G(X_A^r - x_1 - \dots - x_{r-1})g(x_1) \dots g(x_{r-1})dx_1 \dots dx_{r-1}$.

This means that a L_2 bidder A wants to win the auction for sure when $X_A^r \geq E[X_A^r]$ and lose it for sure when $X_A^r < E[X_A^r]$. Therefore his optimal bid coincides with the maximum bid of a L_1 bidder B when $X_A^r \geq E[X_A^r]$ and with the minimum bid of a L_1 bidder B when $X_A^r < E[X_A^r]$. Formally:

$$b_A^r(X_A^r) = \begin{cases} \bar{X}_B^r + E[X_\emptyset^r] + E[X_A^r] & \text{if } X_A^r \geq E[X_A^r] \\ \underline{X}_B^r + E[X_\emptyset^r] + E[X_A^r] & \text{if } X_A^r < E[X_A^r] \end{cases}$$

and the result follows. \square

It is interesting to note that the bid of L_1 is the same as the bid of a fully cursed subject ($\chi = 1$). This happens because the optimal strategy in a second-price auction is to bid the expected value of the good conditional on the information available and the strategy of the opponent. For both L_1 and $\chi = 1$, there is no relationship between the opponent's bid and his information. It is therefore optimal in both cases to bid the unconditional expected value of the rival's private information.

As for the higher levels, remember that L_1 overbids when his private information is below average and underbids when it is above average. L_2 's best response then consists of corner solutions: when L_1 overbids L_2 wants to lose the auction for sure whereas when L_1 underbids L_2 wants to win the auction for sure.

Table 5
Normal Estimation of CE and Lk models.

		All rounds	Round 1	Round 2	Round 3
	χ	1.00	1.00	1.00	1.00
Cursed Equilibrium	AIC	66,917		66,921	
	BIC	66,924		66,942	
	LL	33,458		33,458	
Level-k	L_0	0.22	0.21	0.19	0.15
	L_1	0.75	0.79	0.81	0.84
	L_2	0.00	0.00	0.00	0.01
	Eq	0.03	0.00	0.00	0.00
	AIC	41,571		41,583	
	BIC	41,605		41,686	
	LL	20,780		20,776	

5.3. Estimation

We now estimate these two behavioral models to check how well they each fit the data. For both models we perform two sets of estimations. First, we constrain the parameters to be the same for all rounds, then we allow the parameters to differ across rounds. If the CE and Lk models are robust and capture the reaction of subjects to all three types of information, the parameters in the constrained estimation and in all the rounds of the unconstrained estimations should be similar. Also, since both behavioral models predict NE choices for all parameter values when there is no private information, they cannot be identified. Therefore we leave rounds 0 and 4 out of the estimation.

We use the following econometric specification:

$$b_{om} = b_w(X_{om}) + \varepsilon_{om}$$

where b_{om} is the bid observed in the data for observation o of subject m , $b_w(X_{om})$ is the bid predicted by model $w \in \{CE, Lk\}$, and ε_{om} is an error term assumed to be independently distributed and following the normal distribution $N(0, \sigma)$. Therefore, $\Pr[b_{om} | X_{om}] = f(b_{om} - b_w(X_{om}))$, where $f(\cdot)$ is the density of a normal distribution $N(0, \sigma)$.¹⁴ For each model we find the parameters that maximize the log-likelihood of our sample. Since the Lk model assigns types to subjects, we construct the likelihood function per subject and then sum all subjects' likelihoods. Therefore, if we have O observations per subject, M subjects and L types with proportions π_l that sum to one, we get the following log-likelihood function:

$$LL(\pi, \sigma | b) = \sum_{m=1}^M \log \left(\sum_{l=1}^L \pi_l \prod_{o=1}^O \Pr[b_{om} | X_{om}, l] \right)$$

We assume there is only one type in the CE model and the bidding function has the level of cursedness χ as the only parameter (see later for a brief discussion of heterogeneous cursed players). Therefore, the log-likelihood function is:

$$LL(\chi, \sigma | b) = \sum_{m=1}^M \sum_{o=1}^O \log \left(\Pr[b_{om} | X_{om}, \chi] \right)$$

Table 5 displays the estimation results of the two behavioral models. The column labeled *all rounds* reports the findings when we constrain the parameters to be the same in all rounds. For each of the two sets of estimations, we compute two information criteria: the Akaike information criterion (AIC) and the Bayesian information criterion (BIC). The lower the AIC and BIC values, the better the model fits the data. For both models, the unconstrained estimation is not significantly different than the constrained

¹³ For level-3, there is an interval of possible bids for each value of private information, making it difficult to identify this type. Also, even though truncating the hierarchy might result in not capturing higher levels of sophistication in the data, the inclusion of equilibrium players (Eq , who are more sophisticated than level-3) partially addresses that problem.

¹⁴ An alternative estimation strategy was to use a logit model. However, given the large number of actions, many of which were not observed in the data, we ran into an empty cell problem.

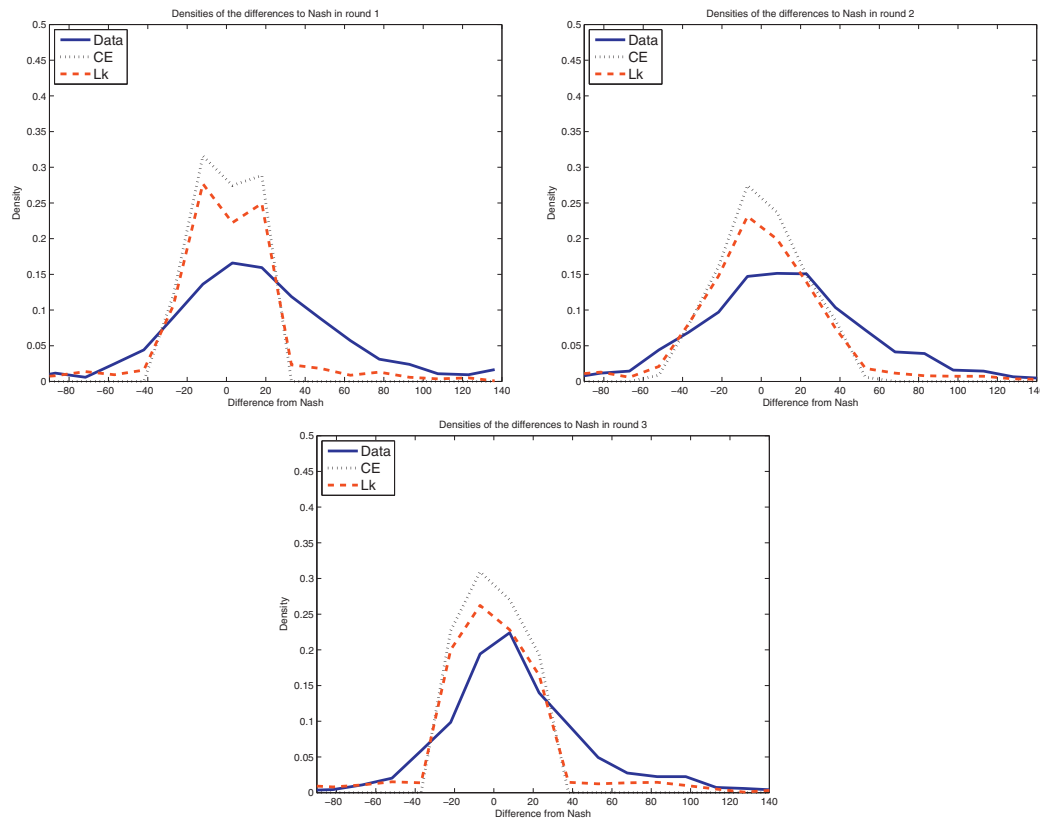


Fig. 4. P.d.f. of deviations from NE in rounds 1, 2 and 3.

one under either criterion, which means that both models have consistent estimates across rounds.

The best estimate of the CE model is $\chi = 1$ in every round, which means that subjects in our sample are fully cursed. This finding follows naturally from the regression analysis of Section 4.3, which showed that subjects react less to private information than predicted by Nash theory and overbid common uncertainty (higher intercept and smaller slope, just like a χ -cursed subject). In fact, subjects react even less to private information than a fully cursed bidder since the slope of the bidding function is less than 1, but no χ parameter can accommodate such severe under-reaction to private information. The observed overbidding cannot be fully captured either since CE predicts an unconditional average bid equal to NE theory and the data shows aggregate overbidding.

The estimates of the Lk model are also in agreement with the estimates of the CE model as 75% of the population or more is L_1 (the equivalent of fully cursed). The rest are L_0 which are meant to capture subjects with very low and very high bids. Lk fits the data better than CE under both the AIC and BIC criteria. This must necessarily be the case in our sample. Indeed, since the best estimate of the (one-parameter) cursedness is $\chi = 1$, the (three-parameter) level-k can do as well by having all subjects as L_1 and weakly better by having a positive fraction of some other types of subjects as well. In that respect, the higher dimensionality of Lk plays strongly in its favor.¹⁵ It is interesting that there are virtually no ‘Equilibrium’ players in the Lk specification. Note also that Lk does not offer a big improvement over CE. This is because our subjects

overbid on average and neither L_0 nor L_2 subjects (let alone E_q) can account for this departure. Finally, Crawford and Iriberry (2007) estimate these two behavioral models with the data of Avery and Kagel (1997). Since our round 2 is equivalent to their setting, we can compare the results. In their constant precision specification (the one we are using in our study), they find $\chi = 0.8$ for the CE model and 94% of L_1 and 6% of L_2 for the Lk model. Our estimations are $\chi = 1$ for CE and 81% of L_1 and 19% of L_0 for Lk, which are similar (strong cursedness in the CE model and overwhelming majority of L_1 players in the Lk model) but not identical.¹⁶

The results are presented in a more intuitive way in Fig. 4. It plots the probability density functions of the deviations from NE observed in the data (solid line) and estimated by the CE (dotted line) and Lk (dashed line) models. On the x-axis are the deviations from NE and on the y-axis are the empirical and estimated probability density functions. The data is less concentrated around the mean and has thicker tails due to a more heterogeneous bidding than predicted by the models. The densities of the two behavioral models have very similar shapes, which is expected given that L_1 and $\chi = 1$ behave identically. The only noticeable difference is the slightly thicker tails of the Lk model due to the presence of L_0 subjects whose bids are, by definition, highly dispersed.

Overall, both models can parsimoniously explain two important features of the data: overbidding when the realized value of private information is small and under-reaction to increases in that information. However, two caveats should be noted. First, neither model can explain the other departures from the theoretical predictions, especially departures with respect to common uncertainty and public information. Indeed, in round 0 (where

¹⁵ This is of course not true in general (e.g. when the best estimate of cursedness is $\chi \in (0, 1)$). Eyster and Rabin (2005) provide an estimation of CE with heterogeneous levels of cursedness. In our experiment this could improve the fit but it is unlikely to do so significantly, as the estimated cursedness is already maximal. It could capture some heterogeneity of the data but it would still fail to account for the aggregate over-bidding.

¹⁶ When Crawford and Iriberry (2007) allow for subject specific precision, the percentage of other types increases. Although this alternative specification has an interesting econometric motivation, we prefer to stick to the more parsimonious behavioral model.

there is no information) subjects bid more than the expected value of the good. They somehow carry this overbidding—on top of the winner's curse—to the rounds with private information and, to a lesser extent, also to round 4 where subjects bid for an item of known value. So, although the theories go in the right direction, our results suggest that there is still room for improvement in the understanding of subjects' reaction to uncertainty and public information, possibly through a combination of behavioral theories. Second, both CE and Lk perform similarly due to the fact that, in our auction, the choice of fully cursed and level-1 players "coincide". However, these identical behaviors are based on widely different decision processes (ignoring the link between other's choices and their information v. assuming a non-strategic choice of the rival and best-responding to it). The current experiment cannot disentangle between these two theories but suggests that a suitably designed variant could potentially determine which process captures best the departures observed in an auction.

The findings in this section are summarized in the following result.

Result 4. Hypothesis 3 is supported by the data: CE and Lk account for a number of important features of the data. However, there is heterogeneity, overbidding and under-reaction to private information beyond the winner's curse that still remain largely unexplained.

6. Conclusion

This and our companion paper (Brocas et al., 2015) have introduced a simple and intuitive design of a common value auction with two distinctive features: we have varied in a parsimonious way the amount of information regarding the value of the good and we have considered three types of information and changed their relative importance across auctions. Taken together, both papers allow us to conclude that a significant fraction of subjects do not fully grasp differences in *types* of information. In particular, they treat private and public information in a similar way. Although more evidence is needed to assess the importance of this behavioral anomaly, the evidence points to a significant departure that deserves further theoretical and empirical scrutiny.

This anomaly has implications in settings beyond common value auctions. Just to name a few examples, juries have access to public information (facts revealed by defense attorney and prosecutor) and private information (prior knowledge of the case and personal interpretation of the law). Consumers are exposed to public messages (advertising) as well as private messages (individual research and pre-purchase tests). Voters obtain news from public sources (newspapers) and private sources (personal understanding of the agenda of a candidate). Our experimental evidence suggests that suboptimal decisions may occur because individuals do not realize that, in strategic contexts, the source of knowledge (public or private) is key in determining the optimal course of action.

Finally, some specific findings of this paper raise other questions that would be interesting to address in future research. First, the fact that subjects react less to public information than what Nash theory predicts is intriguing. It implies for example that subjects in a common value auction would overbid more if the good is drawn from a uniform distribution $[0, 10]$ than if it is drawn from a uniform distribution $[90, 100]$. It would be interesting to check if the under-reaction to public information is an anchoring effect due to our design or if it is indeed a pervasive problem in common value auctions. Second, common value second-price auctions with two bidders have the methodological drawback of featuring multiple equilibria. The literature (including this paper) typically focuses on the symmetric equilibrium. Alternatively, one could study a three-bidder version of this game where the Nash

equilibrium is unique and symmetric, so that the comparison between empirical and theoretical predictions does not rely on a (potentially problematic) symmetry assumption.

Appendix A. Instructions

This is an experiment in group decision making, and you will be paid for your participation in cash at the end of the experiment. Different participants may earn different amounts. What you earn depends partly on your decisions, partly on the decisions of others, and partly on chance. The entire experiment will take place through computer terminals, and all interaction between participants will take place through the computers. It is important that you not talk or in any way try to communicate with other participants during the experiments.

We will start with a brief instruction period. During the instruction period, you will be given a complete description of the experiment and will be shown how to use the computers. You must take a quiz after the instruction period. So it is important that you listen carefully. If you have any questions during the instruction period, raise your hand and your question will be answered so everyone can hear. If any difficulties arise after the experiment has begun, raise your hand, and an experimenter will come and assist you.

At the end of the session, you will be paid the sum of what you have earned in all matches, plus the show-up fee of \$ 5.00. Everyone will be paid in private and you are under no obligation to tell others how much you earned.

Your earnings during the experiment are denominated in tokens. You start the experiment with an endowment of 300 tokens. Depending on your decisions, you can earn more tokens or lose some tokens. At the end of the experiment, we will count the number of tokens you have and you will be paid \$1.00 for every 20 tokens. So you start the experiment with an endowment of 15 dollars.

The experiment will consist of 15 matches. In each match, you will be paired with one of the other participants in the experiment. Since there are participants in today's session, there will be pairs in each match. You are not told the identity of the participant you are matched with. Your payoff depends only on your decisions, the decisions of the one participant you are matched with and on chance. What happens in the other pairs has no effect on your payoff and vice versa. Your decisions are not revealed to participants in the other pairs.

We will now explain how each match proceeds. At the beginning of the match, the computer pairs you with another participant. Next, the computer randomly assigns with equal probability a role to each member as 'bidder 1' or 'bidder 2'. In each match, each member of the pair will be asked to bid for one item. The following screenshot illustrates how the value of the item is calculated.

[SCREEN 1]

You will never see such a screen, but it is useful to understand the game. In this screen, there are 4 boxes [POINT]. Each box contains a number of tokens which is equally likely to be any integer from 0 to 50. The number of tokens in one box is independent of the number of tokens in the other boxes. To be more precise, before each match begins the computer selects, with equal chance, a number of tokens between 0 and 50 for the first box [POINT]; then it selects, with equal chance, a number of tokens between 0 and 50 for the second box [POINT], and so on until the 4th box [POINT]. The value of the item is the sum of the tokens in all four boxes. In this example, the value is 64. Both participants will bid for this value. Note that given the number of tokens in each box varies from 0 to 50, the value of the item is always at least 0 and at most 200. The number of tokens in each box and the corresponding value of the item remain the same during the entire match. However, this information will not be displayed

to you all at once. Instead, the information will be revealed sequentially.

More precisely, each match is divided into 5 rounds. Participants keep the same role and pair for all the rounds of the match. In each round of the match, there is one auction for the whole item. We now explain how the auction in each round works.

Round 0

[SCREEN 2]

Bidder 1 sees a screen similar to the upper part of the slide [POINT]. Bidder 2 sees a screen similar to the lower part of the slide [POINT]. Both bidders see all 4 boxes but they do not see what is inside any of them.

Now you submit a bid for the entire item, that is, for the total number of tokens in all 4 boxes. (We will explain in a minute how bids are transformed into payoffs). You do not get to see the bid of the other participant. Instead, you move to round 1.

Round 1

[SCREEN 3]

In round 1, you keep the same role and bid against the same participant as in round 0. The screens that bidders 1 and 2 see are similar to the upper and lower part of the slide [POINT]. Bidder 1 sees the number of tokens inside the first box starting from the left, the underscoring box [POINT]. Bidder 2 sees the number of tokens inside the first box starting from the right, the overscoring box [POINT]. Note that knowing the number of tokens in one box does not give you any information about the number of tokens in the other boxes. If you are bidder 1, you know that bidder 2 can only observe the content of the overscoring box [POINT] and cannot observe the content of the underscoring box and the boxes that are neither underscoring nor overscoring [POINT]. The analogous reasoning applies to bidder 2.

After observing the content of the underscoring box, if you are bidder 1, or the content of the overscoring box, if you are bidder 2, you submit a bid for the total number of tokens in all 4 boxes. You do not get to see the bid of the other participant. Instead, you move to round 2.

Round 2

[SCREEN 4]

In round 2, you keep the same role and bid against the same participant as in round 1. The screens that bidders 1 and 2 see are similar to the upper and lower part of the slide [POINT]. Bidder 1 now sees the number of tokens inside the first two boxes starting from the left, the two underscoring boxes [POINT]. Bidder 2 sees the number of tokens inside the two boxes starting from the right, the two overscoring boxes [POINT]. Notice that the number of tokens in the leftmost and rightmost boxes did not change [POINT]. This happens because, as we explained before, the tokens inside each box were all drawn at the beginning of the match and do not change between rounds. They are just sequentially revealed to bidders.

After observing the content of the open boxes you submit a bid for the total number of tokens in all 4 boxes. You do not see the bid of the other participant. Instead, you move to round 3.

This process of seeing one more box continues round after round until both bidders have seen the content of all the 4 boxes.

At the end of round 4, when all bids have been made, the computer screen displays the total number of tokens in all boxes. This is the value that both participants were bidding for. Then, the computer randomly selects with equal probability one of the 5 rounds. For the round selected, payoffs are computed as follows. The participant who submitted the highest bid in the selected round wins the total number of tokens in all 4 boxes and pays the bid of the other participant in that round. This payoff can be positive (if the tokens in the boxes exceed the bid of the other participant), zero or negative. This amount is added or subtracted to the current stock of tokens. The participant who submits the

lowest bid pays nothing and obtains nothing; his payoff is zero. The bids in all the other rounds do not count for the payoffs. If both participants submit exactly the same bid, then the computer randomly chooses the winner with equal probability and computes the payoffs just like before.

Remember that in each round you always bid for the tokens in all the boxes, including those for which the content is hidden. You do not have to write the same bid in all rounds of the match. You can increase or decrease your bid from round to round, if you think this will increase your payoff. There are only two restrictions in the bids. First, it has to be an integer number between 0, the minimum value if all the boxes have 0 tokens, and 200, the maximum value if all the boxes have 50 tokens. Second, it cannot exceed your current stock of tokens displayed at the beginning of the match.

When the match is finished, we proceed to the next match. For the next match, the computer randomly reassigns all participants to a new pair, a new role as bidder 1 or bidder 2, and randomly selects the number of tokens to put inside each box. The new assignments do not depend in any way on the past decisions of any participant including you and are done completely randomly by the computer. The assignments are independent across pairs, across participants and across matches. This second match then follows the same rules and payoffs as the first match. Your final payoff in the experiment is equal to your stock of tokens in the end. Basically, it is equal to your initial stock of tokens plus your accumulated payoffs during the experiment.

This continues for 15 matches, after which the experiment ends.

[SCREENS 5 AND 6]

These slides summarize the rules of the experiment.

We will now begin the Practice session and go through one practice match to familiarize you with the computer interface and the procedures. During the practice match, please do not hit any keys until you are asked to, and when you enter information, please do exactly as asked. Remember, you are not paid for this practice match. At the end of the practice match you will have to answer some review questions.

[AUTHENTICATE CLIENTS]

Please double click on the icon on your desktop that says . When the computer prompts you for your name, type your First and Last name. Then click SUBMIT and wait for further instructions.

[START GAME]

[SCREEN 7]

You now see the first screen of the experiment on your computer. It should look similar to this screen.

[Point to overhead screen as you explain this]

At the top left of the screen, you see your subject ID [POINT]. Please record that ID in your record sheet. You have been paired by the computer with one other participant and assigned a role as bidder 1 or bidder 2, which you can see on the top left of the screen [POINT]. The participant you are paired with has been assigned the opposite role (bidder 2 or bidder 1). The pair assignment and role will remain the same for the entire match. You can also see on the top left of the screen that you are in round 0 [POINT].

At this round you cannot see the content of any box.

If you look at the middle of your screen, you should see 'Choose a bid for Round 0' [POINT]. Please type the day of your birthday and press enter. For example, if you were born on March 12, write 12. This is only for the practice match, in the actual experiment you can type any integer number between 0 and 200.

Round 0 is over. We now move to round 1.

[SCREEN 8]

You now see the screen of round 1. It should look similar to this screen. [Point to overhead screen as you explain this]

If you are bidder 1, you can see the content of the underscored box, while if you are bidder 2 you can see the content of the overscored box.

Please type 200 minus the day of your birthday in the field in front of ‘Choose a bid for Round 1’ and press enter.

Round 1 is over. We now move to round 2.

[SCREEN 9]

You now see the screen of round 2. It is the same as the previous rounds, except that both you and the other participant can see the content of one more box [POINT]. The contents of the leftmost and rightmost boxes are the same as in the previous round [POINT]. This happens because the tokens inside each box were all drawn at the beginning of the match.

Please type the month of your birthday in the field in front of ‘Choose a bid for Round 2’ and press enter.

Round 2 is over. We now move to round 3.

[SCREEN 10]

You now see the screen of round 3. It is the same as the previous rounds, except that both you and the other participant can see the content of one more box. Notice that the two boxes in the middle are both underscored and overscored [POINT]. This means that both you and the other participant can observe their content. Please type 100 + the month of your birthday in front of ‘Choose a bid for Round 3’.

Round 3 is over. We now move to round 4.

[SCREEN 11]

You now see the screen of round 4. It is the same as the previous rounds, the only difference is that you can see the content of one more box [POINT]. Please type the first three digits of the year you were born in the field in front of ‘Choose a bid for Round 4’ and press enter.

Round 4 is over. Now the computer will randomly select one the rounds.

[SCREEN 12]

The selected round is highlighted in yellow [POINT]. The participant with the highest bid will collect this amount [POINT] and pay the bid of the other participant. The participant with the lowest bid has a payoff of zero in this match [POINT].

The bottom part of your screen contains a table summarizing the results for all matches you have participated in [POINT]. This is called your history screen. It will be filled out as the experiment proceeds. It only shows the results from your pair, not the results from any of the other pairs. Now click ‘Continue’. The practice match is over. Please complete the quiz. Raise your hand if you have any question.

[WAIT for everyone to finish the Quiz]

Are there any questions before we begin with the paid session? We will now begin with the 15 paid matches. Please pull out your dividers. If there are any problems or questions from this point on, raise your hand and an experimenter will come and assist you.

[START MATCH 1]

[After MATCH 15 read:]

This was the last match of the experiment. Your payoff is displayed on your screen. Please record this payoff in your record sheet and remember to CLICK OK after you are done.

[CLICK ON WRITE OUTPUT]

Your Total Payoff is this amount plus the show-up fee of \$ 5.00. You will pay each of you in private in the next room in the order of your Subject ID number. Remember you are under no obligation to reveal your earnings to the other participants.

Please put the mouse behind the computer and do not use either the mouse or the keyboard. Please remain seated and keep the dividers pulled out until we call you to be paid. Do not converse with the other participants or use your cell phone. Thank you for your cooperation.

Could the person with ID number 0 go to the next room to be paid.

[CALL all the participants in sequence by their ID]

References

- Abbink, K., Brandts, J., 2008. 24. Pricing in bertrand competition with increasing marginal costs. *Games Econ. Behav.* 63, 1–31.
- Albers, W., Harstad, R., 1991. Common-value auctions with independent information: A framing effect observed in a market game. In: Selten, R. (Ed.), *Game Equilibrium Models*, Springer, Berlin, vol. 2, pp. 308–336.
- Andreoni, J., Che, Y., Kim, J., 2007. Asymmetric information about rival's types in standard auctions: an experiment. *Games Econ. Behav.* 59, 240–259.
- Avery, C., Kagel, J.H., 1997. Second-price auctions with asymmetric payoffs: an experimental investigation. *J. Econ. Manag. Strategy* 6 (3), 573–603.
- Brocas, I., Carrillo, J., Castro, M., 2015. The nature of information and its effect on bidding behavior: laboratory evidence in a first-price common value auction. *J. Econ. Behav. Organiz.* 109, 26–40.
- Camerer, C., Ho, T., Chong, J., 2004. A cognitive hierarchy model of behavior in games. *Q. J. Econ.* 119, 861–898.
- Carrillo, J.D., Palfrey, T.R., 2009. The compromise game: two-sided adverse selection in the laboratory. *Am. Econ. J.* 1 (1), 151–181.
- Carrillo, J.D., Palfrey, T.R., 2011. No trade. *Games Econ. Behav.* 71 (1), 66–87.
- Casari, M., Ham, J., Kagel, J., 2007. Selection bias, demographic effects and ability effects in common value auctions experiments. *Am. Econ. Rev.* 97 (4), 1278–1304.
- Charness, G., Levin, D., 2009. The origin of the winner's curse: a laboratory study. *Am. Econ. J.* 1 (1), 207–236.
- Crawford, V.P., Iriberry, N., 2007. 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.
- Eyster, E., Rabin, M., 2005. Cursed equilibrium. *Econometrica* 73 (5), 1623–1672.
- Grosskopf, B., Rentschler, L., Sarin, R., 2010. An Experimental Investigation of Asymmetric Information in Common-Value Auctions. Working Paper.
- Kagel, J.H., Levin, D., 2008. Auctions: A survey of experimental research, 1995–2008. In: Kagel, J.H., Roth, A.E. (Eds.), *The Handbook of Experimental Economics*, vol. II. Princeton University Press.
- Kagel, J.H., Levin, D., Harstad, R.M., 1995. Comparative static effects of number of bidders and public information on behavior in second-price common value auctions. *Int. J. Game Theory* 24 (3), 293–319.
- Klemperer, P., 1998. Auctions with almost common values. *Eur. Econ. Rev.* 42, 757–769.
- Mares, V., Shor, M., 2008. Industry concentration in common value auctions: theory and evidence. *Economic Theory* 35, 37–56.
- Milgrom, P., 1981. Rational expectations, information acquisition, and competitive bidding. *Econometrica* 49 (4), 921–943.
- Milgrom, P.R., Weber, R.J., 1982. A theory of auctions and competitive bidding. *Econometrica* 50 (5), 1089–1122.
- Nagel, R., 1995. Unraveling in guessing games: an experimental study. *Am. Econ. Rev.* 85, 1313–1326.
- Rogers, B., Palfrey, T., Camerer, C., 2009. Heterogeneous quantal response equilibrium and cognitive hierarchies. *J. Econ. Theory* 144 (4), 1440–1467.
- Samuelson, W., Bazerman, M., 1985. The Winner's Curse in Bilateral Negotiations. In: Smith, V. (Ed.), *Research in Experimental Economics*, vol. 3. Greenwich, CT: JAI Press, p. 105–137.
- Stahl, D.O., Wilson, P., 1995. On players' models of other players: theory and experimental evidence. *Games Econ. Behav.* 10, 213–254.