

Bidding in first-price and second-price interdependent-values auctions: A laboratory experiment*

Theodore L. Turocy Timothy N. Cason
University of East Anglia Purdue University

December 18, 2015

Abstract

We report a laboratory experiment on first-price and second-price auctions in settings with independent signals and interdependent values. The environment includes independent private values and the common-value “wallet game” as limiting cases. We manipulate the degree of interdependence of values across sessions, while maintaining the same Bayes-Nash equilibrium bidding function. In contrast, cursed equilibrium predicts bids will be raised for lower signals. We find some support for cursed equilibrium, in that bids change as the degree of value interdependence changes. Contrary to both Bayes-Nash and cursed equilibrium, auction revenues are largest for intermediate levels of interdependence. We construct a model combining cursedness with an underweighting of the opportunity costs of higher bids, and find substantial bidder heterogeneity. A majority of bidders are either fully cursed and disregard completely the bad news that winning the auction entails, or are not cursed at all. We also find evidence for some systematic procedural differences in bidding between first-price and second-price auctions.

Keywords: auctions, affiliated values, winner’s curse, wallet game, experiments.

JEL Classifications: D44, C91.

*The authors thank participants at the 2013 North American Meetings of the Economic Science Association (Santa Cruz), and seminar participants at University of Arizona, Brigham Young University, UC-Santa Barbara, University of East Anglia, University of Edinburgh, University of New South Wales, Monash University, Purdue University, Universidad Torcuato di Tella, and Wissenschaftszentrum Berlin für Sozialforschung (WZB) for helpful comments. Cadence Huen provided valuable assistance in carrying out the experimental sessions through the UEA Springboard internship program. This work was supported by the British Academy through its Small Grants Programme, award number SG113289. Estimations of the quantal response model were carried out on the High Performance Computing Cluster supported by the Research and Specialist Computing Support service at the University of East Anglia.

1 Introduction

In thirty years of laboratory experiments studying auctions for single indivisible objects, most studies have found that bidding tends to exceed the predictions of risk-neutral Nash equilibrium. This overbidding has been documented systematically in both first-price and second-price auctions, as well as in auctions with both pure private and pure common values. The magnitude of the effect, and the resulting loss of earnings by participants relative to the Nash prediction, depends on the environment. Overbidding is typically less prevalent in the second-price private-values auction, although even the presence of a weakly dominant strategy does not eliminate it (Kagel and Levin, 2014). In common-value settings, overbidding is usually large and persistent; even with the salient feedback of winning an auction with a negative profit, many continue to make aggressive bids (Kagel and Levin, 2002).

This paper reports a laboratory experiment designed to look simultaneously at the effects of the pricing rule and the relationship between the information bidders receive and their posterior valuation, within a unified experimental protocol. In two-bidder auction games, we retain the standard assumption that bidders receive private information in the form of statistically independent signals, but vary the extent to which values are interdependent. One motivation for this design is that evidence on bidder behavior has mainly been collected for the private-values and common-values environments in separate experimental studies. Our integrated design includes both of these settings and adds a novel intermediate case of interdependent values. We also consider first-price and second-price rules within the same experimental protocol.

Overbidding relative to the Bayes-Nash equilibrium prediction in first-price auctions with independent private values is a robust phenomenon. Several explanations for this behavior have been advanced. Goeree et al. (2002) provides a thorough treatment of some leading hypotheses using a discretized first-price private-values auction, motivated by the “flat-maximum” critique of Harrison (1989). They manipulate across treatments the costs of over-bidding and under-bidding, relative to the Nash prediction, and they formalize the qualitative predictions of Harrison’s argument using quantal response equilibrium (McKelvey and Palfrey, 1995). They find that the quantal response model combined with constant relative risk aversion, as originally analyzed by Cox et al. (1988), organizes their data well.

Goeree et al. also investigate other competing hypotheses, including fitting a model with a probability weighting function. The best fit model turns out to have a shape that is mathematically equivalent to a risk-aversion formulation, and not the inverted-S shape suggested by e.g. Prelec (1998). This reveals that the menu of lotteries implied by bids in the first-price auction are not well-suited for distinguishing risk aversion and nonlinear probability weighting. A similar idea was pursued by Dorsey and Razzolini (2003), who directly compared participant decisions in first-

price auctions with those in equivalent lottery choice settings. Another competing hypothesis posits a joy of winning; that is, an extra utility for being the winner of the auction. Goeree et al. consider a model with an extra utility for winning, also originally studied by Cox et al. (1988). They find that the joy of winning model does less well in accounting for the data than the model in which bidders are assumed to bid as-if risk averse.

Models based on asymmetric regret are a variant of the joy-of-winning idea which have attracted more attention in recent years. This idea was originally expressed formally by Engelbrecht-Wiggans (1989). This class of models considers outcomes of auctions ex-post. If a bidder wins the auction with a bid strictly higher than his opponents', then after the fact he might regret that he did not bid lower and therefore earn more; this is "winner regret." Conversely, if a bidder loses the auction but the winning bid was below his reservation value, after the fact he might regret he did not bid higher and therefore make positive earnings; this is "loser regret." The hypothesis is that loser regret is more salient than winner regret, and therefore biases bids upwards. Several experimental papers have found evidence consistent with the model, including Filiz-Ozbay and Ozbay (2007), Engelbrecht-Wiggans and Katok (2007, 2008, 2009), and Turocy and Watson (2012). The impulse balance equilibrium of Ockenfels and Selten (2005) incorporates a similar idea.

In contrast, the separate literature on common-value auctions largely attributes overbidding to the hypothesis that bidders systematically fail to understand the negative informational consequences of winning the auction, and therefore fall prey to the "winner's curse." Kagel and Levin (2002) present an extensive survey and discussion of the experimental evidence. An important formalization of this failure of updating is the cursed equilibrium of Eyster and Rabin (2005), in which players may only partially, and therefore improperly, take into account the fact that the behavior of other bidders depends on their private information.

A qualitative prediction of cursed equilibrium in some settings, including the one studied here, is that bids for low signal draws should be disproportionately aggressive. Avery and Kagel (1997) conduct experiments with two-bidder second-price auctions, in which the value of the object is the sum of uniform independent signals received privately by each bidder. They find that expected profits for bidders with low signals are persistently negative, both within a session, and even with "experienced" bidders. In first-price auctions, Holt and Sherman (2000) conduct a small number of classroom-type experiments, where the value of the object is the average of uniform independent signals.¹ They, too, find evidence that matches the qualitative features of cursed equilibrium.

The environment in the present paper encompasses both independent private values and the pure-common-value wallet game, while admitting a range of intermediate cases in which bidders have differing values, which depend on the realization of both bidders' signals. Investigation of af-

¹Both settings are therefore equivalent to the "wallet game" auction which is a staple classroom demonstration; see e.g. Klemperer (1998).

filiated value environments other than the well-studied pure-common and completely-independent values environments is important considering that most auctions in practice have both common- and private-value components.²

Our setting retains independent and uniformly-distributed private signals. The symmetric Bayes-Nash equilibrium does not depend on the extent of interdependence between bidders' values, and is therefore the same as the well-studied case of pure private values. Retaining uniform independent private signals allows for a protocol that changes as little in the instructions as possible, and maintains the same bidding and feedback interface across different valuation structures. By keeping the baseline Bayes-Nash equilibrium the same, we can isolate the extent to which bidding in interdependent-values and common-values cases is driven by the cursed equilibrium's (mis)perception of the informational content of winning the auction, as opposed to biases or heuristics which may be applicable also in the private-values case.

We estimated a structural model of bidding to determine the distribution of cursed belief updating. We find substantial heterogeneity across bidders in terms of cursedness. When values are interdependent, many bidders appear to completely disregard the bad news that winning the auction entails (i.e., they are "fully" cursed), while many others fully account for it. More than half of our participants fall into one of these two boundary cases, illustrating that measures of "average" degrees of cursedness need not represent the typical participant's behavior.

We also find that auction revenues are largest for an intermediate level of interdependence of values, which neither Nash nor cursed equilibrium alone predicts. Our bid function estimates account for this via a combination of cursedness and an underweighting of the opportunity costs of higher bids. We also report evidence that suggests some systematic procedural differences in bidding in first-price and second-price auctions.

In comparing private- and common-values environments through the lens of cursed equilibrium, our design is similar to a recent paper by Cox (2015), who considers contribution behavior in a threshold public goods game. Echoing our results, he finds mixed evidence, with cursed equilibrium organizing some but not all qualitative deviations of behavior from the Bayes-Nash equilibrium prediction. Our study also shares some of the research objectives pursued by Crawford and Iriberry (2007), who also seek to understand overbidding in common- and private-value settings and consider cursed equilibrium. Their main interest is in initial strategic thinking and not learning or equilibrium, however, so their empirical analysis focuses on initial bids submitted by inexperienced subjects.

²Only Forsythe et al. (1989), Goeree and Offerman (2002) and Kirchkamp and Moldovanu (2004) have previously considered interdependent but not pure-common-value structures in experiments. They implement interdependence in different ways, and focus on other research issues such as the revenue and efficiency impacts of additional information disclosure. In contrast to the present experiment, none of these studies include both first- and second-price auction rules or a comparison with pure private- or common-value structures.

The paper is organized as follows. Section 2 formally states the class of auction environments we consider, and derives predictions for the Bayes-Nash and cursed equilibria in the setting. Section 3 outlines the experimental design and protocol. Section 4 presents the data and results. Section 5 concludes with a discussion.

2 Theory

There are $N + 1$ bidders, and each bidder $i = 1, \dots, N + 1$ receives an independently drawn signal x_i from the uniform distribution on $[0, 1]$. The value of the object to bidder i is equal to a weighted average of his signal, and the maximum of the other signals:³

$$u_i(x_1, \dots, x_{N+1}) = (1 - \gamma)x_i + \gamma \max_{j \neq i} x_j. \quad (1)$$

The case $\gamma = 0$ reduces to the standard independent private-values auction. When there are two bidders and $\gamma = \frac{1}{2}$, this is a re-scaled version of the “wallet game.”

We analyze the equilibrium of this auction, under both first-price and second-price rules, using the cursed equilibrium of Eyster and Rabin (2005), which includes the standard Bayes-Nash equilibrium as a special case. In this context cursed equilibrium captures the idea that bidders may not fully understand the informational implications of winning the auction, which is a prominent type of failure in Bayesian updating that is broadly consistent with the observed persistence of the winner’s curse in common-value auctions. A cursed equilibrium is characterized by a parameter $\chi \in [0, 1]$, which captures the degree of cursedness; that is, the degree to which a bidder underestimates the informational impact of winning the auction. If $\chi = 0$, bidders are good Bayesians; if $\chi = 1$, they are “fully cursed.” A χ -cursed bidder bids as if his posterior expected value, conditional on winning, equals $(1 - \chi)$ times the “correct” posterior expectation, and χ times a “naive” posterior expectation, in which he takes account of his own signal, but only the (unconditional) expectation of the other bidders’ signals.

Eyster and Rabin point out that the formulation of the cursed equilibrium fits within the framework of Milgrom and Weber (1982), so the symmetric cursed equilibrium of both first-price and second-price auctions follow directly from the theorems therein. For the first-price auction, the χ -cursed equilibrium is

$$b^*(x) = \frac{N}{N + 1}(1 - \gamma\chi)x + \frac{N}{N + 1}\gamma\chi, \quad (2)$$

³In introducing this value structure, Turocy (2008) suggests that it captures in a simple way the possibility of later resale; if the winning bidder has to re-sell the object at a later date, then the amount he is likely to be able to sell it is proportional to the highest signal received by other bidders.

and for the second-price auction it is

$$b^*(x) = (1 - \gamma\chi)x + \frac{1}{2}\gamma\chi. \quad (3)$$

This implies the following predictions about both the first-price and second-price auctions:

1. The standard Bayes-Nash equilibrium (i.e., when $\chi = 0$) bidding function is independent of γ .
2. When $\gamma > 0$ (i.e., the auction is not pure-private-values), a χ -cursed bidder adopts a bid function with a positive bid at the lowest signal $x = 0$, and a lower slope relative to the Bayes-Nash equilibrium prediction.

In the first-price auction the bid submitted at the highest signal $x = 1$ is always $b^*(1) = \frac{N}{N+1}$, independent of γ or χ . The bid function rotates up and flattens with increases in either γ or χ . In the second-price auction, cursed equilibrium generates a rotation of the bidding function around the midpoint of the range of signals; it predicts an upward bias in bids for low signals, but a downward bias for high signals. Therefore, specific shifts in bidding behavior provide evidence for cursed equilibrium beliefs.

3 Experimental design

We report 24 experimental sessions conducted at the University of East Anglia, each using 8 undergraduate students drawn from a subject pool maintained via ORSEE (Greiner, 2015) by the Centre for Behavioural and Experimental Social Science. None of the 192 participants took part in more than one session. The full text of the instructions is available in a supplementary appendix. Participants knew that there would be 40 periods, and that they would be re-paired into two-bidder auctions in each period. All interaction among participants was fully computerized, and they had no access to participant IDs or other identifying information regarding either their current co-player or the co-players in their history.

We employ a 3×2 factorial design, varying (between sessions) in one dimension the degree of interdependence of values, and in the other whether the first-price (FPA) or second-price (SPA) payment rule was used. For the interdependence of values, we considered three cases: $\gamma = 0$, the case of pure private values (abbreviated PV below); $\gamma = \frac{1}{4}$, a case of interdependent affiliated values (IV); and $\gamma = \frac{1}{2}$, which is a re-scaled version of the pure common-values wallet game (CV).

Participants each period bid in a sealed-bid auction. Each received a draw of a private signal from the set $\{0.20, 0.40, \dots, 9.80, 10.00\}$.⁴ To maximize comparability across sessions and

⁴All signals, bids, and earnings were expressed directly in pounds sterling; we avoided the use of in-lab currency

treatments, the same realizations of private signals were used in all sessions, as well as the same sequence of participant ID pairings.

At the beginning of each auction period, subjects' computer screens displayed their signal and a 5-second countdown clock. After the countdown participants could select a bid using a custom slider device which was used previously in private-values auctions in Turocy et al. (2007) and Turocy and Watson (2012). Bids could be submitted in increments of 0.10.⁵ Bids were specified by clicking a point along the slider, and then confirming with a button click. The bidding period lasted at least 40 seconds, or until 5 seconds after the last bidder had submitted her bid. The results of the auctions were displayed for 15 seconds, after which the next period began.⁶

In addition to providing feedback in the same graphical frame as used by the participant to set a bid, the feedback screen also provided information on the results of the auction from the perspective of the other bidder. Figure 1 shows sample screenshots showing a typical outcome in the first-price auction in the private-values and common-values settings, respectively. We chose this rich graphical presentation for two reasons. First, we wanted to ensure that the process of generating valuations from signals was as transparent as possible, without requiring participants to follow along with the arithmetic in detail. Second, given the stylized facts that losses are a persistent feature in common-value auctions, we designed an environment in which a participant could learn not only from the experience of her own losses, but also potentially from the negative consequences of aggressive bidding by others in the session.

[Figure 1 about here.]

Participants also had on their screen a record sheet showing the complete history of all auctions they had participated in. This record sheet was displayed at all times, including while the subject was bidding, while waiting after bidding for the period to complete, and during the feedback interval between periods.

To maintain approximate equivalence of earnings in the baseline risk-neutral equilibrium, participants in CV sessions received an initial balance of £8.00, and in IV sessions an initial balance of £4.00. Ten out of the 40 periods were selected at the end of the session for payment. This allows a string of negative earnings in early periods to not necessarily cause bankruptcy, eliminating a

units and associated exchange rates. The exchange rate between GBP and USD at the time of the experiments was approximately 1 GBP = 1.60 USD.

⁵We selected the granularity of the signal and bid spaces to be fine enough that ties in either signals or bids would be relatively unlikely, making the continuous approximation reasonable. The discretization of signals in increments of 0.20 and bids in increments of 0.10 ensures that bidding one-half the signal is always possible, and remains an equilibrium of the first-price auction in the discretized game.

⁶We designed this pacing of the auction periods to control for the subjects' opportunity cost of time, as no subject could make the session conclude faster by bidding faster, as well as to make the private signal and the feedback process salient. Few bidders took more than 40 seconds to submit a bid in any period.

design complication which arises in many common-value experiments. The initial balances were set large enough that the chances of negative total earnings were small, even under assumptions of very aggressive bidding.⁷ Average earnings were £12.16, with a standard deviation of £5.50, and an interquartile range of [£8.18, £15.75]. Sessions lasted about 65 to 70 minutes, inclusive of instructions and payment.

4 Results

4.1 Overview

As a first view of the data, Table 1 presents the aggregate summary statistics for each of the six treatments, averaging over all periods of all sessions.⁸

[Table 1 about here.]

Qualitatively the data match with several received stylized facts. Bidding is quite aggressive in the first-price private-values auction, leading to revenues far in excess of the risk-neutral Nash prediction. Bidding in the second-price private-values auction results in revenues slightly higher than the dominant-strategy prediction; as we will see below, this is driven by a core of dominant-strategy bidders combined with varying levels of overbidding by some participants. A novel result is that for both pricing rules, revenues are greatest in the interdependent-values auctions. Both private-values settings result in a high proportion of auctions (over 90%) being won by the bidder with the highest signal. We interpret this as a measure of the homogeneity of bidding across bidders, as this measure would be 100% if all bidders adopted the same pure-strategy bidding function (whether equilibrium or not).⁹ We note that the result that risk attitudes should be irrelevant in FPA-CV (Holt and Sherman, 2000) but not in FPA-PV implies, other things being equal, a higher proportion of auctions being won by the high-signal bidder in FPA-CV, as heterogeneity in risk attitudes would not drive heterogeneity in bidding. Our data indicate exactly the opposite trend.

4.2 Individual bids

[Figure 2 about here.]

⁷Two of the 192 subjects, both in the IV treatment, had negative earnings after the 10 paying periods were drawn; these subjects received a payment of zero.

⁸In expectation, both the Nash and cursed equilibrium revenues are the same for the first-price and second-price auctions. The slightly higher predictions for the second-price auction are an artefact of the particular realization of signals.

⁹For PV and IV, the high signal win percentage is equivalent to the proportion of efficient allocations. For CV, efficiency does not depend on which bidder receives the object.

[Figure 3 about here.]

We next consider the full joint distribution of bids and signals. Figure 2 (for FPA) and Figure 3 (for SPA) provide two views on this relationship. To the left is a scatterplot of all bids.¹⁰ These are complemented to the right with panels displaying the median bid and interquartile range of bids for each decile of signals. These figures show substantial variation of individual bids, deviating considerably from the linear theoretical predictions. In all treatments, however, the median bid is roughly an affine function of the signal. Bids as a function of signal are similar between FPA-PV and FPA-IV, other than the bids above signal in FPA-IV for low signals. For high signals, bids appear more conservative in the FPA-CV. In the second-price auction, there is some qualitative evidence of the rotation in bid functions predicted by cursed equilibrium for CV; however, this is not (as) evident in the IV treatments.

4.3 Consistency with dominant strategy bidding

In treatment SPA-PV, it is a weakly dominant strategy to bid equal to the signal. It is therefore a useful check on the procedures and bidding interface to consider the frequency with which bidding conforms to the weakly dominant strategy. The top panels of Figure 3 already show considerable clustering around signal bidding in this treatment. For a more formal comparison we focus on the second half of the session periods to exclude early, noisy behavior reflecting initial learning. We look both at whether bids overall are consistent with the weakly dominant strategy, and whether individual bidders typically follow the strategy. We classify a bid as being consistent with the weakly dominant strategy if it is equal to the signal, or one bid increment (£0.10) below.¹¹ We classify a bidder as being a dominant strategy bidder if at least 75% of their bids in the second half of the experiment (15 out of 20 periods) are consistent with the weakly dominant strategy.

[Table 2 about here.]

Table 2 presents a summary of dominant strategy bids in this treatment. Overall, slightly more than 50% of bids are consistent with the weakly dominant strategy, and 19 of 36 bidders are classified as consistently dominant strategy bidders. This rate of dominant strategy bidding is substantially higher than most of the literature on experimental second-price PV auctions. For example, the recent study of Agranov and Yariv (2015) also features two bidders and random re-pairing each period, and a large majority of the observed bids exceed value. As summarized in the survey by Kagel and Levin (2014), the frequency of dominant strategy bidding in second-price

¹⁰Points in the scatterplots are jittered to avoid having the same (signal, bid) pairs exactly overlay each other.

¹¹We found several bidders consistently chose to bid one increment below. Given the discretization of our signal and bid space, if all other bidders bid equal to their signal, this would still be a best response.

PV auctions more typically ranges between 20 and 30%, even when bidders have considerable field experience in auctions (Garratt et al., 2012). The novel feedback structure in our graphical interface may be one reason for this greater frequency of dominant strategy play.

4.4 Earnings left on the table

As bids tend to be above the risk-neutral Nash prediction, participants earn less than Nash equilibrium predicts. To get a measure of the strength of the financial costs suffered by bidders due to their aggressive bidding, we construct a measure of bidders’ optimization premiums against the signals and bids they actually faced, or could have faced, in their session.

Formally, let $x_{i,t}$ denote the signal received by bidder i in period t , and $b_{i,t}$ the bid she submitted. Let $T(t) = \lfloor t/10 \rfloor$ partition the 40 periods of the session into four phases of 10 periods each. For each bidder i , we take each realized signal and bid, $(x_{i,t}, b_{i,t})$, and compute the bidder’s expected payoff by simulating its performance against each of the 70 realized signal-bid combinations of other bidders in that same phase in that session. We write this set of possible competing signal-bid combinations as¹²

$$C_{i,t} = \{(x_{j,s}, b_{j,s}) : j \neq i \wedge T(t) = T(s)\}.$$

Then, for that same $x_{i,t}$, we determine the best-response bid $b_{i,t}^*(x_{i,t})$ treating the set $C_{i,t}$ as an estimator of the joint distribution of potential opponents’ signals and bids. The difference between the expected earnings of the best response $b_{i,t}^*(x_{i,t})$ and the earnings of the actual bid $b_{i,t}$ provides a measure of the money “left on the table” due to suboptimal bidding, against the actual behavior of others in the same session.

[Table 3 about here.]

Table 3 presents the average amounts left on the table for each treatment. As measured by foregone earnings, bidding in FPA is farther from optimal than SPA. Per ten periods, average foregone earnings ranged from £3.16 to £4.82 in the early periods, which is substantial when compared to the average actual take-away earnings of £12.16 across all sessions. Over time, in FPA there is a clear trend of bids moving towards the (risk-neutral) best response, with money left on the table being almost halved from the first phase of 10 periods to the last. The trend in SPA is less clear. Only in PV is a trend towards more optimal bidding evident, which is attributable to some bidders realizing the dominant strategy as the session progresses.

¹²We measure against all signal and bid realizations instead of only the ones actually faced by the bidder, because, in principle, a bidder could have been matched against any other bidder in a given period. This calculation mirrors the perspective a bidder has at the point in time when she has learned her signal but not the signal and bid of her rival.

4.5 Session-level results

We next break down the data by independent session to begin some formal hypothesis testing across treatments. Table 4 presents statistics on each of the 24 sessions, grouped by treatment. In this table, we introduce two additional measures of the aggressiveness of bidding. The excess bid for a bidder in a period is his actual bid minus the risk-neutral Nash equilibrium prediction; the excess revenue for an auction is the actual revenue minus the risk-neutral Nash prediction. For each session, we report the median of these measures over all periods.

[Table 4 about here.]

Holding the auction rule fixed, at the standard Bayes-Nash equilibrium these measures are all equivalent across the three value structures. In cursed equilibrium, however, revenues increase when moving from PV to IV to CV, and for the first-price auction bids also increase when moving from PV to IV to CV. Our first formal result provides little support for these predictions.

Result 1. *Mean revenues, median excess bids and median excess revenues are higher in FPA-IV than FPA-CV; median excess bids and revenues are higher in FPA-PV than FPA-CV; and median excess bids are higher in SPA-IV than both SPA-PV and SPA-CV.*

Support. We conduct nonparametric Mann-Whitney-Wilcoxon tests comparing the four measures across value structures, holding fixed the auction rule, using independent sessions as the unit of observation (4 observations per treatment). For these tests we use only data from the second half of the sessions (i.e., the last 20 periods) to focus on behavior after the initial learning phase. The pairwise comparisons summarized in the result are significant at the 5% level (two-tailed tests). □

Although there is little systematic effect of the value structure on bids and revenues, the choice of the pricing rule has a clear effect. Consistent with and extending the stylized facts well-documented in the independent private-value environment, revenue equivalence can be rejected.

Result 2. *Median excess revenues are higher in the first-price auction than the second-price auction for all three value structures.*

Support. Table 4 shows, for each value structure, that the smallest median excess revenue in any first-price auction session is greater than the largest median excess revenue in any second-price auction session. The non-overlapping distributions for these independent session observations imply that the Mann-Whitney-Wilcoxon test rejects the null hypothesis of equal excess revenues, against the two-sided alternative, at the 5% level. □

Cursed equilibrium implies greater differences in bidding behavior for the lowest signal draws, and for the highest signal draws in the second-price auction. In particular, for the lowest draws bids should *increase* when moving from PV to IV to CV; and for the highest draws in the second-price auction the bids should *decrease* when moving from PV to IV to CV. Our next result focuses on the lowest and highest 20% of draws to test this model implication for bidding, using the ratio of bid/signal draw to allow for comparisons across the range of draws.

Result 3. *Bids submitted for low signal draws are lowest in the PV environment for the first-price auction, and bids submitted for high signal draws are highest in the IV environment for the second-price auction.*

Support. Table 5 presents the median bid/signal ratio for the low and high draws, averaged across sessions. Mann-Whitney-Wilcoxon tests indicate that for the FPA, the mean ratio for the PV environment (0.67) for low signals is significantly lower than for the other two value environments at the 5% level (two-tailed tests). The IV and CV ratios are not significantly different, although the overall ordering for these low signal draws is consistent with the cursed equilibrium prediction of more overbidding for more common values.¹³ For the SPA, none of the ratios are significantly different for the low signal draws. For the high signal draws, however, the mean ratio for the IV environment (1.04) is significantly greater than for the other two value environments (again two-tailed Mann-Whitney-Wilcoxon tests at the 5% level.) □

[Table 5 about here.]

The general pattern is that bidding is most aggressive in the IV treatment. This result is novel, as no previous experiments have directly compared interdependent-values environments to pure private- and common-value environments. Importantly, note that this more aggressive bidding for interdependent values is not consistent with the cursed bidding model for a constant χ , which predicts a monotonic relationship as one moves from PV through IV to CV for the first-price auction.

4.6 Estimation of bid functions

Participants in the experiment are heterogeneous in their approaches to bidding, both in terms of their bid levels and in the variability of their bids from period to period. To capture these features of the dataset, this section reports summary information on structural estimates of individual bid functions for all bidders.¹⁴

¹³We omit the ratios for the high signal draws for the FPA because they are expected to be similar across the value environments in both the cursed and standard Bayes-Nash equilibrium models.

¹⁴The online supplementary appendix to this paper, available at <http://www.gambit-project.org/turocy>, provides bidder-by-bidder plots of the fitted bid functions, as well as full data and source code for the estimation.

We define the decision utility function of a bidder as depending on two parameters, $\chi \in [0, 1]$ and $\alpha \in [0, \infty)$. In IV and CV environments, bidders may have χ -cursed valuations, as in Eyster and Rabin (2005). While χ -cursedness generates bids in the first-price auction that are above the risk-neutral Nash prediction when any value interdependence exists, it cannot provide an account of such aggressive bidding in the FPA-PV case. The second parameter α allows the model to account in a flexible way for incentives to submit bids above or below the risk-neutral prediction, as we explain next.

In FPA, the comparison between a bid b and the next higher bid $b + \delta$ involves a tradeoff between two considerations. The higher bid $b + \delta$ wins in some contingencies in which b would lose. A bidder (presumably) sets this against the additional cost of δ the bidder will pay in those contingencies in which b would also win the auction. Note that any quasi-concave decision utility function that predicts bids above risk-neutral in FPA-PV must, at the risk-neutral optimal bid, value the incremental gain more highly than the incremental cost.

Consider the bidding problem from Bidder 1's perspective. Let x denote Bidder 1's signal, and ξ the signal of Bidder 2, which at the point of bidding is a random variable from Bidder 1's perspective. Let $v(x, \xi)$ be Bidder 1's value for the object, conditional on the signals being x and ξ . The probability of the joint event that Bidder 2's signal is ξ and he bids β is $\pi(\beta, \xi)$. The expected earnings for Bidder 1 from a given bid b , conditional on the signal x , are

$$E^{FPA}(b|x) = \frac{1}{2} \sum_{\xi} (v(x, \xi) - b)\pi(b, \xi) + \sum_{\xi} \sum_{\beta < b} (v(x, \xi) - b)\pi(\beta, \xi).$$

The first term represents tied bids (resolved with a fair coin flip in our experiment).

Let δ be the minimum bid increment, noting that $\delta = 0.10$ in our experiment. Then, the difference in expected earnings between bids b and $b + \delta$ is

$$\begin{aligned} E^{FPA}(b + \delta|x) - E^{FPA}(b|x) &= \frac{1}{2} \sum_{\xi} (v(x, \xi) - (b + \delta))\pi(b + \delta, \xi) + \frac{1}{2} \sum_{\xi} (v(x, \xi) - b)\pi(b, \xi) \\ &\quad - \delta \sum_{\xi} \sum_{\beta < b + \delta} \pi(\beta, \xi). \end{aligned} \quad (4)$$

The first term captures the case that $b + \delta$ leads to a tie, which Bidder 1 wins; the second term the case that bidding b would lead to a tie which Bidder 1 would lose. The final term covers the cases in which both b and $b + \delta$ win with certainty; in this case $b + \delta$ results in an additional payment of δ .

We can decompose the terms in (4) into two parts, distinguished by whether the difference in expected earnings from higher bids results in explicit benefits and costs or implicit opportunity costs. As in most auction experiments, our instructions refer to $v(x, \xi)$ as a "resale value," and

explain that earnings are the difference between the resale value and the price paid. The first two terms in (4) therefore refer to the earnings that a participant would realize by bidding $b + \delta$ instead of b , in the relevant contingencies. We distinguish the costs in the third term in (4) by whether the bid is above or below his value. If the bidder would have won with a bid b below his value, then the incremental loss in bidding $b + \delta$ instead of b is an opportunity cost, in the sense that the bidder could have earned more in those contingencies by bidding b . If $b + \delta$ exceeds the value of the object, then the extra δ contributes to an explicit loss. Let $\mathbf{1}(p)$ be the indicator function taking on values 1 when the logical predicate p is true, and 0 when it p is false. We then write the decomposition of (4) for non-cursed bidders as $DB^{FPA,E} - DC^{FPA,E}$ where

$$\begin{aligned}
DB^{FPA,E}(b + \delta|x) &= \frac{1}{2} \sum_{\xi} (v(x, \xi) - (b + \delta))\pi(b + \delta, \xi) + \frac{1}{2} \sum_{\xi} (v(x, \xi) - b)\pi(b, \xi) \\
&\quad - \delta \sum_{\xi} \sum_{\beta < b + \delta} \mathbf{1}(v(x, \xi) < b + \delta)\pi(\beta, \xi) \\
DC^{FPA,E}(b + \delta|x) &= \delta \sum_{\xi} \sum_{\beta < b + \delta} \mathbf{1}(v(x, \xi) \geq b + \delta)\pi(\beta, \xi).
\end{aligned}$$

For cursed bidders, let $\bar{\xi}$ be the unconditional expected value of the signal of Bidder 2. The decomposition of (4) is then $DB^{FPA,C} - DC^{FPA,C}$ where

$$\begin{aligned}
DB^{FPA,C}(b + \delta|x) &= \frac{1}{2} \sum_{\xi} (v(x, \bar{\xi}) - (b + \delta))\pi(b + \delta, \xi) + \frac{1}{2} \sum_{\xi} (v(x, \bar{\xi}) - b)\pi(b, \xi) \\
&\quad - \delta \sum_{\xi} \sum_{\beta < b + \delta} \mathbf{1}(v(x, \bar{\xi}) < b + \delta)\pi(\beta, \xi) \\
DC^{FPA,C}(b + \delta|x) &= \delta \sum_{\xi} \sum_{\beta < b + \delta} \mathbf{1}(v(x, \bar{\xi}) \geq b + \delta)\pi(\beta, \xi).
\end{aligned}$$

The decision utility gives possibly differing weights to $DB^{FPA,\cdot}$ and $DC^{FPA,\cdot}$. The parameter $\alpha > 0$ captures the weighting of the incremental opportunity cost of a higher bid relative to the incremental realized gains or losses. For given χ and α , we define the decision utility function for FPA as

$$\begin{aligned}
U^{FPA}(b + \delta|x) &= \sum_{\beta \leq b + \delta} [(1 - \chi)DB^{FPA,E}(\beta|x) + \chi DB^{FPA,C}(\beta|x)] \\
&\quad - \alpha \sum_{\beta \leq b + \delta} [(1 - \chi)DC^{FPA,E}(\beta|x) + \chi DC^{FPA,C}(\beta|x)] \quad (5)
\end{aligned}$$

If $\alpha = 1$ no bias exists in the relative weighting of incremental realized and opportunity costs of higher bids, and the model is equivalent to risk neutrality. If $\alpha < 1$, the opportunity costs are

underweighted, and the maximizing bid is greater than the risk-neutral maximizer.

We show in Appendix A that, in the case of bids which do not risk losing money, the model is equivalent to the weighted anticipated regret model of Engelbrecht-Wiggans (1989) for FPA-PV. We also consider in the Appendix the constant relative risk aversion model, which has often appeared in analyses of bidding in FPA-PV. In the Appendix we illustrate that bidding data in FPA-PV which would generate an estimate of a CRRA parameter r (where $r = 0$ is risk neutrality) will correspond roughly to an estimate of $\alpha = 1 - r$. However, as already shown by Holt and Sherman (2000), risk attitudes are irrelevant in FPA-CV, and therefore cannot account for aggressive bids observed there.¹⁵

In SPA, a bidder's own bid never determines his payment, and therefore there is no equivalent to the opportunity cost component of the FPA model. We therefore define decision utility in SPA only to incorporate the possibility of cursedness.¹⁶ The expected payoff to Bidder 1 bidding b , conditional on receiving signal x and recalling that Bidder 2 bids β , is therefore

$$U^{SPA,E}(b|x) = \sum_{\xi} \sum_{\beta < b} (v(x, \xi) - \beta) \pi(\beta, \xi) + \frac{1}{2} \sum_{\xi} (v(x, \xi) - b) \pi(b, \xi)$$

For a cursed bidder, we have

$$U^{SPA,C}(b|x) = \sum_{\xi} \sum_{\beta < b} (v(x, \bar{\xi}) - \beta) \pi(\beta, \xi) + \frac{1}{2} \sum_{\xi} (v(x, \bar{\xi}) - b) \pi(b, \xi)$$

Combining these we then define the SPA decision utility as

$$U^{SPA}(b|x) = (1 - \chi) U^{SPA,E}(b|x) + \chi U^{SPA,C}(b|x). \quad (6)$$

The decision utility function of a bidder is completely specified by the parameters $\chi \in [0, 1]$, which determines his degree of cursedness, and, for FPA, $\alpha \in [0, \infty)$, which determines the weighting of the opportunity cost of higher bids. To take the model to the data, we assume each bidder uses a behavioral strategy in which the logarithm of the probability of choosing each bid b is proportional to the decision utility of b . That is, they “better-respond” vis-a-vis the decision utility rather than best-respond. Each bidder i is characterized by parameters χ_i and α_i , as well as the constant of proportionality $\lambda_i \in [0, \infty)$. Under this logit choice assumption the probability that

¹⁵Also, the possibility of losses is present in FPA-CV and FPA-IV, which cause technical problems in applying CRRA or other models based on risk aversion.

¹⁶In SPA-IV and SPA-CV, the logit choice model we use naturally generates a prediction that bids above the risk-neutral prediction will occur more often than those below. This is because higher bids lead to a smaller reduction in expected payoffs than lower bids.

bidder i chooses a bid of b when he has signal x is given by

$$\Pr(b|x; \chi_i, \alpha_i) = \frac{\exp \lambda_i U(b|x; \chi_i, \alpha_i)}{\sum_{\beta} \exp \lambda_i U(\beta|x; \chi_i, \alpha_i)}. \quad (7)$$

Camerer et al. (2014) use a similar combination of logit choice and cursedness to organize high bids in an independent signals second-price auction in which the value of the object is the maximum of the two signals. We differ from their approach in that they adopt the heterogeneous quantal response equilibrium framework of Rogers et al. (2009); we instead only assume that bidders have correct rational-expectations beliefs about the behavior of others within their own session. For each participant i we take the realized signals and bids of the other participants in the same session, and using that data we compute the expected decision utility $\hat{U}(b|x; \chi, \alpha)$ for each possible signal x and bid b , conditional on the decision utility parameters.¹⁷

We then compute parameter estimates for each participant i , $(\hat{\lambda}_i, \hat{\chi}_i, \hat{\alpha}_i)$, via maximum likelihood. The parameters χ and α have the straightforward interpretations described above. On the other hand, in logit choice the parameter λ is measured in the same units as the decision utility function. The utility scale changes as χ and α change, so we convert these parameters into characteristics of the estimated behavior strategy, motivated by the useful maximum entropy property of the logit choice rule (7). Fix the signal x , and the decision utilities $U(b|x; \chi, \alpha)$ associated with each bid b , and consider the distribution in (7) for some λ . Let $U^*(\lambda)$ denote the expected decision utility participant i will receive if he chooses the distribution of bids from (7). Then, the distribution of bids is the most random distribution, as measured by entropy, which yields at least a payoff of $U^*(\lambda)$.

Therefore, for each participant i , instead of reporting λ_i , we compute the entropy of the joint distribution of signals and bids conditional on $(\hat{\lambda}_i, \hat{\chi}_i, \hat{\alpha}_i)$. This entropy is necessarily bounded below by the entropy of the behavior strategy in which the bidder randomizes uniformly over all bids for each signal, and bounded above by the entropy of a pure behavior strategy. We use this to compute a normalized entropy H_i , with $H_i = 0$ corresponding to uniform randomization and $H_i = 1$ to a pure strategy. Because these are measures of the randomization involved in the estimated behavior strategy, H_i is more comparable than λ_i across participants and treatments.

Table 6 summarizes of the distributions of parameter estimates for five of the treatments. Figure 4 provides histogram plots of those distributions. We omit analysis of SPA-PV, in view of the dominant-strategy nature of the equilibrium and the large number of players who play that dominant strategy in a majority of periods (Table 2). For each treatment, we estimate parameters using

¹⁷We therefore take advantage of the experiment structure. First, we only use the bids from the participant's session; therefore, if there are session effects, they are accounted for directly. Because of the random matching and salient feedback of others' behavior, we believe that a rational-expectations model is appropriate, in that participants had ample opportunity to assess the behavioral patterns of other bidders in their session.

all periods (top panels), as well as conducting separate estimations using only data from the first 20 or last 20 periods as shown in lower panels.

[Table 6 about here.]

[Figure 4 about here.]

Result 4. *The degree of belief cursedness is highly heterogeneous and bimodal. Cursedness alone is not enough to explain aggressive bids in FPA-IV and FPA-CV; cost underweighting is observed in all FPA treatments, and is most pronounced in the IV environment.*

Support. The mean estimate for the cursedness parameter χ is similar across treatments, with means based on all periods ranging from 0.357 in SPA-IV to 0.530 in FPA-CV. However, the bimodal distribution of the individual estimates shown in Figure 4 is striking. The most common parameter estimates for each treatment are fully cursed ($\hat{\chi} = 1$) or fully un-cursed ($\hat{\chi} = 0$). A substantial fraction of participants bid as if they understand, at least qualitatively, the bad news that winning the auction entails when values are common or interdependent ($\hat{\chi} = 0$). Another large fraction bid as if they do not understand this at all ($\hat{\chi} = 1$). \square

Because cursedness is comparable across treatments, and in particular is if anything slightly lower in IV than CV, it cannot alone provide an account for the higher revenues observed in IV than CV. Further, the estimated normalized entropy values \hat{H} are comparable across all FPA treatments and across SPA treatments, so gross differences in the randomness of behavior cannot explain this result either. Estimates of the mean cost-weighting parameter $\hat{\alpha}$ are similar in FPA-PV ($\hat{\alpha} = 0.537$) and FPA-CV ($\hat{\alpha} = 0.607$) and somewhat lower in FPA-IV ($\hat{\alpha} = 0.352$). The FPA-PV estimate is in line with other experiments which use CRRA to explain bidding behavior, where typically a CRRA parameter of around $r = 0.5$ is obtained.

Result 5. *Learning over the course of the session in FPA is accounted for by increases in the opportunity cost weighting parameter $\hat{\alpha}$ rather than changes in cursedness $\hat{\chi}$.*

Support. In all of FPA-PV, FPA-IV, and FPA-CV, the estimated distribution of opportunity cost weights α shifts upwards in the last 20 periods (lowest panels of Table 6 and Figure 4) relative to the first. For each treatment, the mean as well as the three quartile points of the distribution are higher in the last 20 periods. \square

These estimates provide an account which relates the persistent overbidding in FPA in both private and common values environments. Most bidders behave as if they underweight the opportunity cost of winning with an unnecessarily high bid. The degree to which they do so is dispersed across the population, and in all value structures bidders do learn incrementally to take

this opportunity cost more into account. In the IV and CV cases, a proportion of the population of bidders bids as-if fully cursed, and these bidders do not learn to update correctly over time. To a first approximation, cursedness is discrete: bidders either are or are not cursed, and those who are cursed in general do not experience an “a-ha” moment allowing them to transition to uncursed bidding.¹⁸ The lowest distribution of estimates for the cost-weighting parameter $\hat{\alpha}$ in the FPA-IV also contribute significantly to the greatest overbidding observed for the interdependent-value environment.

4.7 Response times

The structure of the bidding period ensured that a period would take at least a minute irrespective of the speed of bidding. This should lower the opportunity cost of contemplating the most appropriate bid to submit each period. Our software tracked usage patterns of the bidding interface, as well as the time each bidder took to confirm the bid in each period. While many factors affect response times, there are interesting patterns in the response time data which provide some exploratory evidence on how participants engage with the bidding task.

[Figure 5 about here.]

Result 6. *In the first-price auction, bidders take longer to bid when they have a high signal, whereas in the second-price auction response times are roughly independent of signal. Bidding is fastest overall in PV-SPA, in which there is a weakly dominant strategy.*

Support. Figure 5 plots the median response time in each treatment, by signal decile. The qualitative features of response times as a function of the signal differ between first-price and second-price auctions. In the first-price auction, response times tend to be longer for higher signals, whereas for the second-price auction, response times are roughly independent of signals.

[Table 7 about here.]

We quantify these differences by estimating a simple regression model with response time as the dependent variable, and the signal, a time trend, and several demographics indicator variables as independent variables.¹⁹ We use a random effect at the participant level, and cluster standard errors at the session level. Table 7 provides the resulting parameter estimates and robust standard

¹⁸We also investigated whether the demographic characteristics used as regressors in Table 7 correlated with bid function parameters; we found no significant relationships. In particular, our male and female bidders exhibit similar distributions of $\hat{\alpha}$ and $\hat{\chi}$. This suggests that our results on the distributions of parameters will not be sensitive to characteristics of the subject pool.

¹⁹We also conducted but do not report a version where the dependent variable is the log of the response time. The pattern of statistical significance is the same in those specifications.

errors. We find that the visual impressions of the effect of signal on response time in the first-price auction are indeed statistically significant.

The positive estimate on the reciprocal of the auction period ($1/t$) indicates that the pace of bidding accelerates over time in the second-price auction, but not generally in the first-price auction, except in IV. This would be consistent, for example, with the hypothesis that the second-price pricing rule is less transparent initially due to the more indirect relationship between bids and earnings; after a few periods of experience with the rule, participants better understand the mechanics and formulate their bids more rapidly. \square

We hypothesize that the effect of signal on response time is driven by the incentives given by the pricing rule. In the first-price auction, a bidder with a high signal is quite likely to win the auction. In that event, her earnings will be determined by the bid she chooses, and as such a bidder might think quite carefully about the bid. On the other hand, in a second-price auction, while the high-signal bidder is still quite likely to win, conditional on winning the bidder's earnings are independent of the bid chosen. Therefore, fine-grained reasoning about bid choice might seem less salient to a bidder even with a high signal.

Point estimates for the effect of the signal in the second-price auction, while statistically insignificant, are negative. In the graphical interface, the confirm bid button was located below the bidding device; therefore, response times for bidders submitting low bids would naturally be smaller and could account for the negative sign. If true, this would imply that the estimates for effect of signal in the first-price auction are understated slightly.

We are not aware of other experiments in auction settings which have systematically considered the effects of the pricing rule or valuation structure on response times. The extensive survey of Spiliopoulos and Ortmann (2014) on response time analysis in experimental economics does not enumerate any auction studies, and the recent focus in economics is on the correlation between response time on errors and social preferences (e.g. Recalde et al., 2015).

5 Conclusion

This experimental design compares pricing rules and valuation structures within a unified experimental protocol, which allows it to identify more cleanly how these dimensions affect bidding behavior. Indeed, while bids remain more aggressive than the risk-neutral Bayes-Nash prediction in all treatments, patterns of behavior differ systematically across treatments and participants respond to both pricing rules and valuation structures. If one were to view this exercise as a competition between Bayes-Nash equilibrium and cursed equilibrium as the most accurate model to approximate auction behavior, we see our results as a split decision. The data provide support

for a key qualitative prediction of cursed equilibrium, that bids should be less responsive to signals when the interdependence of values is greater. Many individual subjects bid as if fully cursed, with estimates of the cursedness parameter χ equal to one. However, others do not have cursed beliefs, and the most aggressive bidding on average is observed in the intermediate case of interdependent but not fully common values. This result cannot be accommodated by either model, although it may arise through a combination of some cursed beliefs and an underweighting of the opportunity cost of incrementally higher bids in FPA.

We used a rich visual interface for providing feedback to participants on the outcome of each auction. This device is able to communicate the relationship between private signals and the value of the object without needing to rely solely on verbal or mathematical descriptions. This visual interface also shows the outcome of the auction period from the perspective of the rival bidder using the same format. We find evidence that the ability to learn from bad outcomes of the other bidder is a function of the pricing rule; bidders become more conservative after seeing the other bidder lose money in the second-price auction, but not in the first-price auction.

The framework within which we conducted this experiment is a step in the direction of understanding the behavioral determinants of bidding in auction experiments. Viewed from the perspective of equation (1), the IV environment, with $\gamma = \frac{1}{4}$, lies between PV ($\gamma = 0$) and CV ($\gamma = \frac{1}{2}$). However, behavior in IV is not bracketed in such a tidy way by the behavior observed in PV and CV. A model combining cursedness and underweighting of bid costs can rationalize this qualitatively, but requires that bid cost underweighting is more extreme in IV. Thus, this is only a partial account.

Although the design aimed to minimize the presentation differences among PV, IV, and CV, and the graphical device avoids reliance on participants needing to carry out the calculation in (1), nevertheless it is plausible that the valuation structure in IV is more cognitively demanding. Models of maximizing, or approximately-maximizing, bidding behavior may not be able to provide a complete account of the relationship between the demands of reasoning about the valuation structure in a game vis-a-vis reasoning about the bid formulation process. Relatedly, some of our results point to procedural differences in bid formulation that are a function of the pricing rule. We find that the pricing rule affects response times, both in terms of how the bidder's signal affects bid times (it does in FPA but not SPA), and in whether bidding speeds up over time (it does in SPA but not in FPA).

These results suggest that our understanding of bid formulation in auctions will be enhanced if we complement models of as-if (approximately) maximizing bidding with a more systematic evaluation of process-related evidence. In our design, the graphical presentation plus control over the length of bidding periods has yielded hints of procedural differences driven by the interaction between pricing rule and valuation structures, and between explanation of the rules of the game

and feedback processing. This graphical approach is amenable not only to the capture of response times, but other measures relevant to process, such as eye-tracking and mouse motion tracking. A more complete account of what participants are doing when bidding in auction experiments would inform not only the analysis of why bidding is aggressive in both private- and common-values settings, but also the evaluation of the external validity of laboratory experiments for understanding bidding processes in the field.

A Notes on the decision utility function

Our decision utility function (5) generalizes the anticipated regret formulation of Engelbrecht-Wiggans (1989) to interdependent valuation structures. From the perspective of Bidder 1, winner regret is defined as the expected difference between Bidder 1's bid, and the actual bid submitted by Bidder 2, when Bidder 1 is the winning bidder. That is,

$$R_W(b|x) = \sum_{\xi} \sum_{\beta < b} (b - \beta) \pi(\beta, \xi).$$

The change in winner regret comparing a bid b to $b + \delta$ is

$$\begin{aligned} R_W(b + \delta|x) - R_W(b|x) &= \sum_{\xi} \sum_{\beta < b + \delta} (b + \delta - \beta) \pi(\beta, \xi) - \sum_{\xi} \sum_{\beta < b} (b - \beta) \pi(\beta, \xi) \\ &= \delta \sum_{\xi} \sum_{\beta < b + \delta} \pi(\beta, \xi), \end{aligned}$$

that is, the change in winner regret is exactly $DC^{FPA,E}(b|x)$ and $DC^{FPA,C}(b|x)$, as previously defined, for the case of bids which are less than the private value. Loser regret is defined as the expected difference between Bidder 1's value and Bidder 2's bid, in the case when Bidder 1 loses the auction, but Bidder 2's bid is below Bidder 1's value. That is,

$$R_L(b|x) = \sum_{\xi} \sum_{\beta > b} (v(t, \xi) - \beta)^+ \pi(\beta, \xi) + \frac{1}{2} \sum_{\xi} (v(t, \xi) - b)^+ \pi(b, \xi),$$

where $(m)^+ \equiv \max\{m, 0\}$. The change in loser regret between b and $b + \delta$ is

$$\begin{aligned}
R_L(b + \delta|x) - R_L(b|x) &= \sum_{\xi} \sum_{\beta > b + \delta} (v(x, \xi) - \beta)^+ \pi(\beta, \xi) - \sum_{\xi} \sum_{\beta > b} (v(x, \xi) - \beta)^+ \pi(\beta, \xi) \\
&\quad + \frac{1}{2} \sum_{\xi} (v(x, \xi) - (b + \delta))^+ \pi(b + \delta, \xi) - \frac{1}{2} \sum_{\xi} (v(x, \xi) - b)^+ \pi(b, \xi) \\
&= - \sum_{\xi} (v(x, \xi) - (b + \delta))^+ \pi(b + \delta, \xi) \\
&\quad + \frac{1}{2} \sum_{\xi} (v(x, \xi) - (b + \delta))^+ \pi(b + \delta, \xi) - \frac{1}{2} \sum_{\xi} (v(x, \xi) - b)^+ \pi(b, \xi) \\
&= - \frac{1}{2} \sum_{\xi} (v(x, \xi) - (b + \delta))^+ \pi(b + \delta, \xi) - \frac{1}{2} \sum_{\xi} (v(x, \xi) - b)^+ \pi(b, \xi).
\end{aligned}$$

For the case of bids strictly less than value, the change in loser regret is exactly $DB^{FPA,E}(b|x)$.

It follows that in the case of $\chi = 0$ and $\alpha = 1$, (5) is just expected monetary earnings, up to a constant which depends on x , and therefore (5) produces the same logit choice distributions as would be obtained by using expected monetary earnings. Similarly, for private values and bids below value, the model corresponds exactly with the anticipated regret formulation of Engelbrecht-Wiggans (1989).

We also note that the standard CRRA utility function $u(m) = m^{1-r}$ often used to model aggressive bidding in FPA-PV likewise implies a similar weighting. It is straightforward to show that, in the case of continuous bids, the maximizing bid b satisfies the first-order condition

$$(x - b)^{-r} \left[(x - b) \frac{d}{db} \Pr(b > \tilde{\beta}) - (1 - r) \Pr(b > \tilde{\beta}) \right] = 0. \quad (8)$$

Therefore, at the maximizing bid, the implied weighting on marginal cost is $\alpha = 1 - r$.

References

- Marina Agranov and Leeat Yariv. Collusion through communication in auctions. Manuscript, 2015.
- Christopher Avery and John H. Kagel. Second-price auctions with asymmetric payoffs: An experimental investigation. *Journal of Economics & Management Strategy*, 6:573–603, 1997.
- Colin Camerer, Salvatore Nunnari, and Thomas R. Palfrey. Quantal response and nonequilibrium beliefs explain overbidding in maximum-value auctions. Manuscript, 2014.
- Caleb Cox. Cursed beliefs with common-value public goods. *Journal of Public Economics*, 121:52–65, 2015.
- James C. Cox, Vernon L. Smith, and James M. Walker. Theory and individual behavior of first-price auctions. *Journal of Risk and Uncertainty*, 1:61–99, 1988.
- Vince P. Crawford and Nagore Iriberri. Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner’s curse and overbidding in private-value auctions? *Econometrica*, 75:1721–1770, 2007.
- Robert Dorsey and Laura Razzolini. Explaining overbidding in first price auctions using controlled lotteries. *Experimental Economics*, 6:123–140, 2003.
- Richard Engelbrecht-Wiggans. The effect of regret on optimal bidding in auctions. *Management Science*, 35:685–692, 1989.
- Richard Engelbrecht-Wiggans and Elena Katok. Regret in auctions: Theory and evidence. *Economic Theory*, 33:81–101, 2007.
- Richard Engelbrecht-Wiggans and Elena Katok. Regret and feedback information in first-price sealed-bid auctions. *Management Science*, 54:808–819, 2008.
- Richard Engelbrecht-Wiggans and Elena Katok. A direct test of risk aversion and regret in first price sealed-bid auctions. *Decision Analysis*, 6:75–86, 2009.
- Erik Eyster and Matthew Rabin. Cursed equilibrium. *Econometrica*, 73:1623–1672, 2005.
- Emel Filiz-Ozbay and Erkut Ozbay. Auctions with anticipated regret: Theory and evidence. *American Economic Review*, 94:1407–1418, 2007.
- Robert Forsythe, R. Mark Isaac, and Thomas R. Palfrey. Theories and tests of ‘blind bidding’ in sealed-bid auctions. *RAND Journal of Economics*, 20:214–238, 1989.
- Rodney J. Garratt, Mark Walker, and John Wooders. Behavior in second-price auctions by highly experienced eBay buyers and sellers. *Experimental Economics*, 15:44–57, 2012.

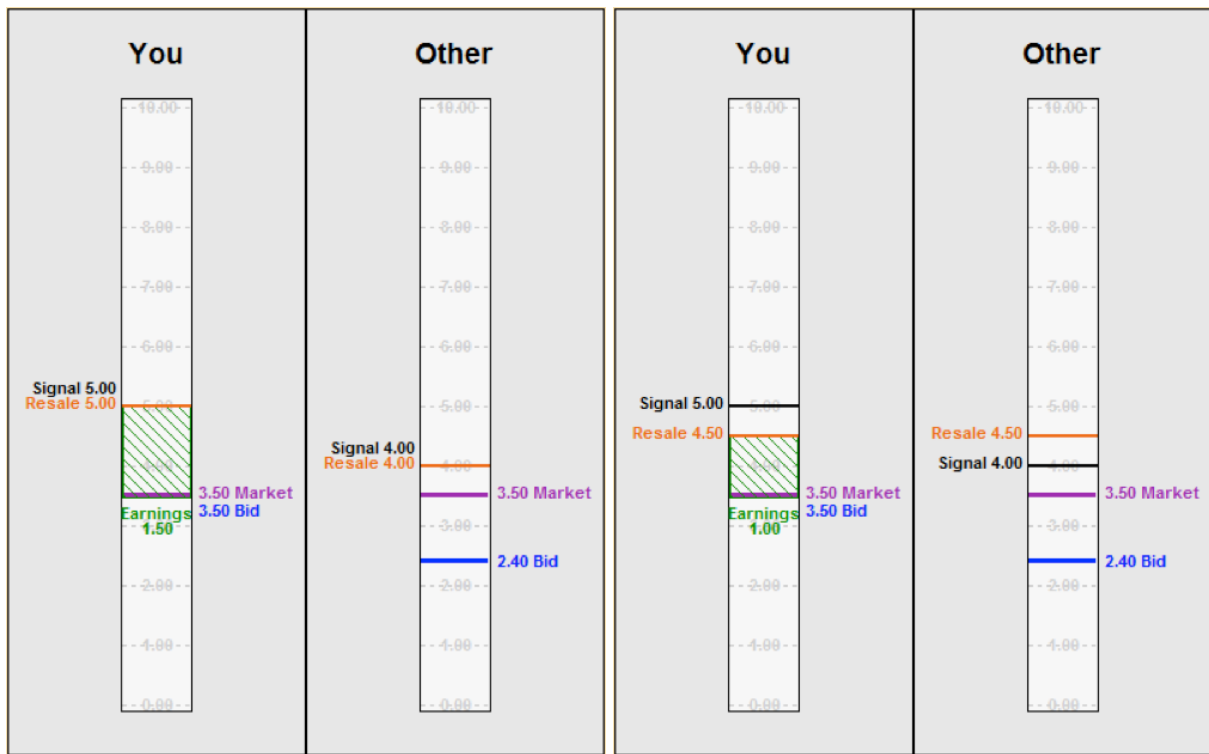
- Jacob K. Goeree and Theo Offerman. Efficiency in auctions with private and common values: An experimental study. *American Economic Review*, 92:625–643, 2002.
- Jacob K. Goeree, Charles A. Holt, and Thomas R. Palfrey. Quantal response equilibrium and overbidding in private-value auctions. *Journal of Economic Theory*, 104:247–272, 2002.
- Ben Greiner. Subject pool recruitment procedures: Organizing experiments with ORSEE. *Journal of the Economic Science Association*, 1:114–125, 2015.
- Glenn W. Harrison. Theory and misbehavior of first-price auctions. *American Economic Review*, 79:749–762, 1989.
- Charles A. Holt and Roger Sherman. Risk aversion and the winner’s curse. Working paper, 2000.
- John H. Kagel and Dan Levin. *Common Value Auctions and the Winner’s Curse*. Princeton University Press, Princeton NJ, 2002.
- John H. Kagel and Dan Levin. Auctions: A survey of experimental research, 1995-2010. In John H. Kagel and Alvin Roth, editors, *The Handbook of Experimental Economics, Volume 2*. 2014.
- Oliver Kirchkamp and Benny Moldovanu. An experimental analysis of auctions with interdependent valuations. *Games and Economic Behavior*, 48:54–85, 2004.
- Paul Klemperer. Auctions with almost common values: The ‘wallet game’ and its applications. *European Economic Review*, 42:757–769, 1998.
- Richard D. McKelvey and Thomas R. Palfrey. Quantal response equilibrium for normal form games. *Games and Economic Behavior*, 10:6–38, 1995.
- Paul R. Milgrom and Robert J. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50:1089–1122, 1982.
- Axel Ockenfels and Reinhard Selten. Impulse balance equilibrium and feedback in first price auctions. *Games and Economic Behavior*, 51:155–170, 2005.
- Drazen Prelec. The probability weighting function. *Econometrica*, 66:497–527, 1998.
- Maria Recalde, Arno Riedl, and Lise Vesterlund. Error prone inference from response time: The case of intuitive generosity in public-good games. University of Pittsburgh Department of Economics WP 15-004, 2015.
- Brian Rogers, Thomas R. Palfrey, and Colin Camerer. Heterogeneous quantal response equilibrium and cognitive hierarchies. *Journal of Economic Theory*, 144:1440–1467, 2009.

Leonidas Spiliopoulos and Andreas Ortmann. The BCD of response time analysis in experimental economics. SSRN Working Paper 2401325, 2014.

Theodore L. Turocy. Auction choice for uncertainty-averse sellers facing strategic uncertainty. *Games and Economic Behavior*, 62:155–179, 2008.

Theodore L. Turocy and Elizabeth Watson. Reservation values and regret in laboratory first-price auctions: Context and bidding behavior. *Southern Economic Journal*, 78:1163–1180, 2012.

Theodore L. Turocy, Elizabeth Watson, and Raymond C. Battalio. Framing the first-price auction. *Experimental Economics*, 10:37–51, 2007.



(a) A private-values auction outcome

(b) A common-values auction outcome

Figure 1: Structure of feedback for a typical auction period. This compares the feedback of the same combination of signals and bids, for private-values and common-values, in a first-price auction.

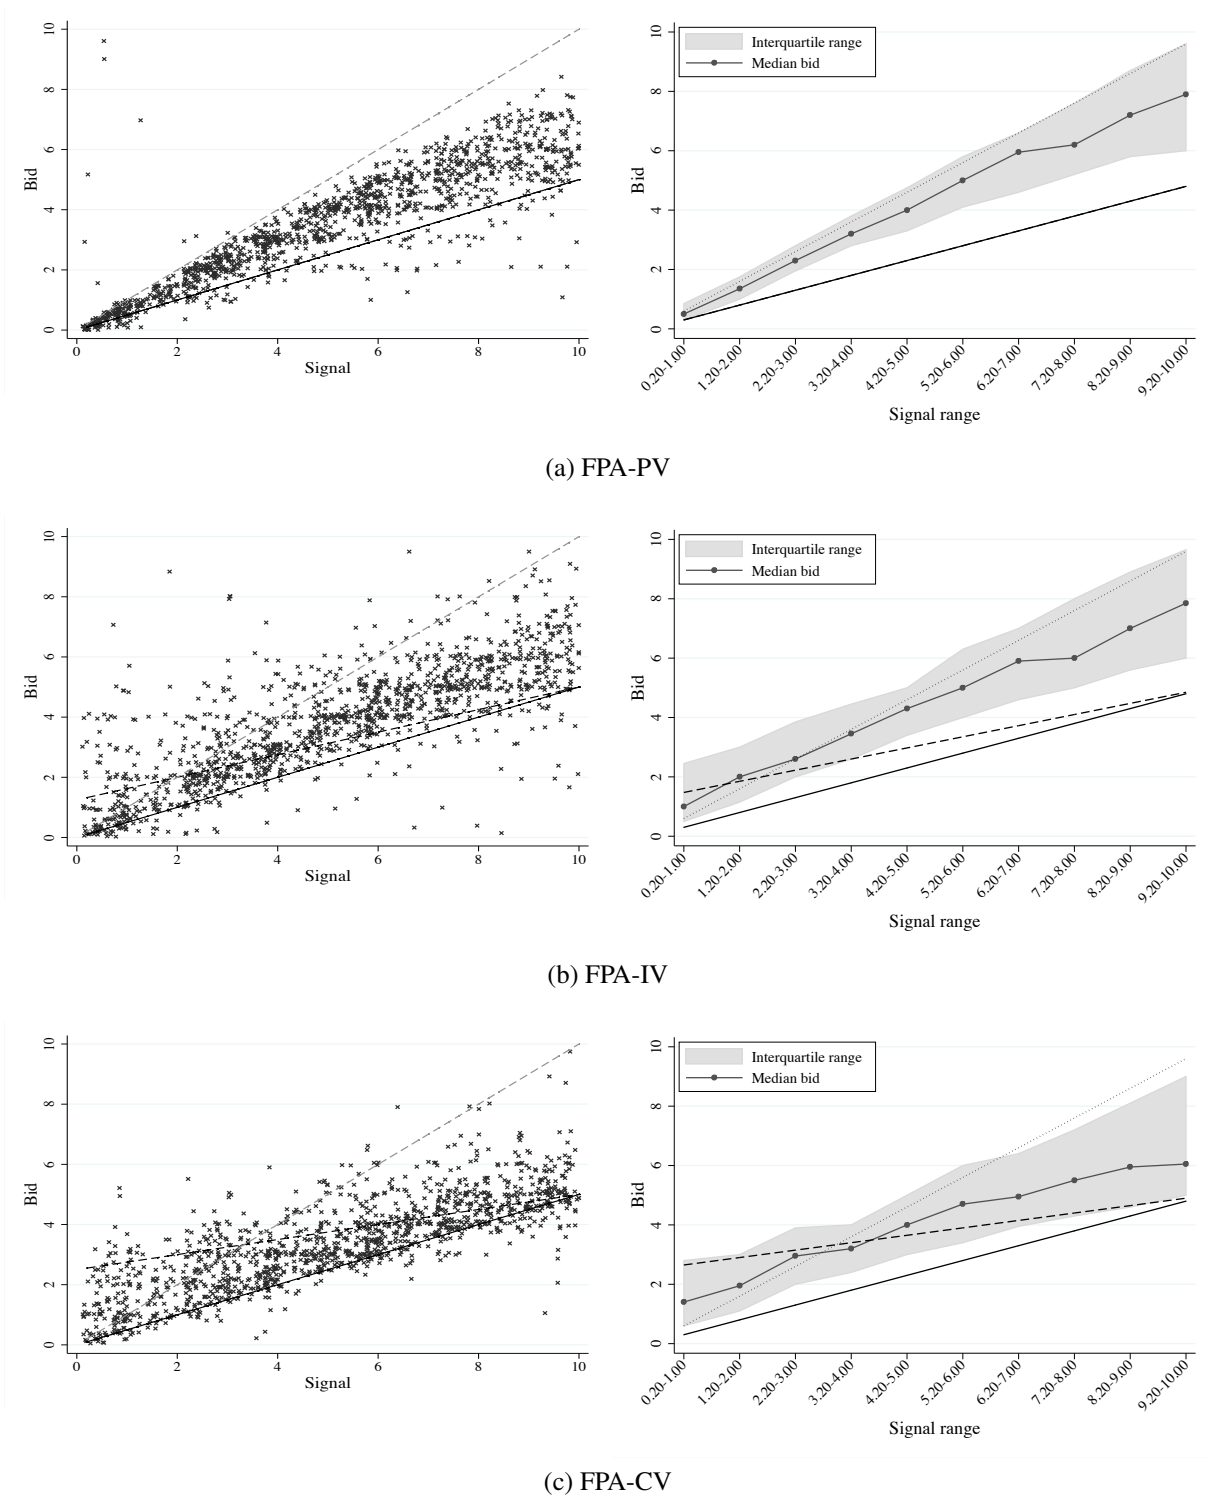
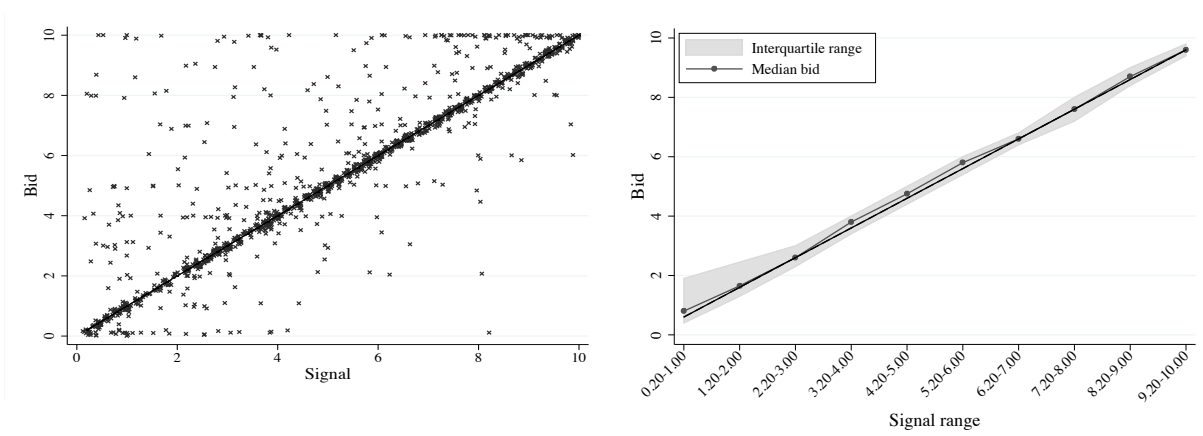




Figure 2: Bids as a function of signal for first-price auctions. For each treatment, the left panel is the scatterplot of all bids. The right panel groups signals by decile, and presents the median and interquartile range of bids for each signal decile. The solid lines correspond to the risk-neutral Nash equilibrium, and the dashed to the fully-cursed equilibrium.



(a) SPA-PV




(b) SPA-IV




(c) SPA-CV


Figure 3: Bids as a function of signal for second-price auctions. For each treatment, the left panel is the scatterplot of all bids. The right panel groups signals by decile, and presents the median and interquartile range of bids for each signal decile. The solid lines correspond to the risk-neutral Nash equilibrium, and the dashed to the fully-cursed equilibrium.



(a) All periods

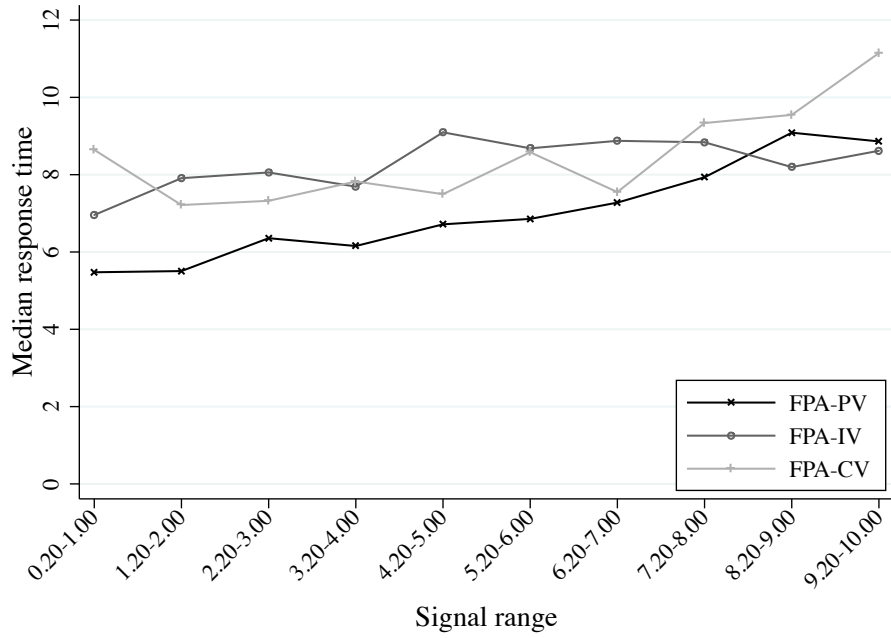


(b) First 20 periods only




(c) Last 20 periods only

Figure 4: Distribution of bid function parameter estimates.



(a) First-price auction



(b) Second-price auction

Figure 5: Median response times by signal decile, all treatments.

Variable	First-Price (FPA)			Second-Price (SPA)		
	PV	IV	CV	PV	IV	CV
Sessions	4	4	4	4	4	4
Mean revenue	4.66	4.92	4.29	3.75	4.12	3.94
(Standard error of mean)	(0.06)	(0.06)	(0.05)	(0.09)	(0.09)	(0.08)
Nash equilibrium revenue	3.40	3.40	3.40	3.47	3.47	3.47
Fully cursed equilibrium revenue	3.40	3.80	4.20	3.47	3.85	4.23
High signal win %	90.6	84.1	83.1	90.3	84.2	81.7
Response time (sec)	9.48	10.45	10.80	6.43	10.38	8.54

Table 1: Summary statistics by treatment, pooling all sessions.

Values	Auction	Session	% dominant strategy	# dominant strategy
PV	SPA	20131115A	0.36	2
PV	SPA	20131127B	0.80	7
PV	SPA	20131127C	0.68	4
PV	SPA	20131202B	0.84	6

Table 2: Percentage of bids at dominant strategy or one increment below, second-price private-values sessions, periods 21-40.

Periods	First-Price (FPA)			Second-Price (SPA)		
	PV	IV	CV	PV	IV	CV
1-10	34.5	48.2	31.6	19.2	16.0	19.1
11-20	29.6	41.6	20.9	13.5	17.1	11.4
21-30	25.4	34.2	20.7	11.4	14.4	15.7
31-40	20.9	25.3	19.8	5.4	14.4	12.6

Table 3: Average money “left on the table” due to suboptimal bidding, by treatment, in pence per period. Values are the difference between the expected payoffs from each bidder’s actual bids in the 10 period block, versus the best-response bids, measured against all realized signals and bids by other bidders in that period block.

Values	Auction	Session	Mean Revenue	High signal Win %	Median excess Bid	Median excess Revenue
PV	FPA	20131016A	4.23	89.4	0.70	0.80
PV	FPA	20131018B	5.07	88.7	1.30	1.65
PV	FPA	20131113A	4.41	90.6	0.85	1.10
PV	FPA	20131113B	4.93	93.8	1.15	1.50
IV	FPA	20131120A	4.89	83.8	1.10	1.35
IV	FPA	20131127A	4.68	89.4	0.90	1.20
IV	FPA	20131202A	4.93	82.4	1.20	1.40
IV	FPA	20131205C	5.18	80.6	1.30	1.80
CV	FPA	20131017A	4.18	84.4	0.35	0.60
CV	FPA	20131017B	4.42	84.4	0.90	1.00
CV	FPA	20131018A	4.05	84.4	0.50	0.50
CV	FPA	20131113C	4.52	79.4	0.80	0.90
PV	SPA	20131115A	4.45	83.1	0.00	0.20
PV	SPA	20131127B	3.71	88.1	0.00	0.00
PV	SPA	20131127C	3.58	92.5	0.00	0.00
PV	SPA	20131202B	3.27	97.5	0.00	0.00
IV	SPA	20131129A	4.12	78.1	0.40	0.20
IV	SPA	20131129B	4.35	88.8	0.40	0.60
IV	SPA	20131205A	4.12	87.5	0.50	0.40
IV	SPA	20131205B	3.90	82.5	0.35	0.30
CV	SPA	20131115B	4.23	75.0	0.20	0.20
CV	SPA	20131118A	3.67	81.9	-0.10	0.00
CV	SPA	20131118B	3.91	81.9	0.00	0.45
CV	SPA	20131202C	3.96	88.1	0.00	0.20

Table 4: Summary of sessions. Measures are computed using all periods. Median excess bid and median excess revenue are measured relative to the risk-neutral Nash equilibrium prediction.

Variable	First-Price (FPA)			Second-Price (SPA)		
	PV	IV	CV	PV	IV	CV
Sessions	4	4	4	4	4	4
Median bid/signal for signals ≤ 2 (Standard error of mean)	0.67 (0.12)	1.02 (0.19)	1.45 (0.26)	1.48 (0.48)	1.89 (0.27)	1.82 (0.26)
Median bid/signal for signals ≥ 8 (Standard error of mean)				1.00 (0.03)	1.04 (0.01)	0.97 (0.03)

Table 5: Mean bid/signal ratios across sessions for low and high signal ranges (periods 21-40).

Estimates using all periods						
	FPA-PV		FPA-IV		FPA-CV	
	Mean	Quartiles	Mean	Quartiles	Mean	Quartiles
\hat{H}	0.323	(0.281, 0.347, 0.386)	0.244	(0.189, 0.251, 0.315)	0.265	(0.205, 0.265, 0.325)
$\hat{\chi}$	—	—	0.531	(0.000, 0.637, 1.000)	0.530	(0.052, 0.592, 0.893)
$\hat{\alpha}$	0.537	(0.331, 0.509, 0.632)	0.352	(0.279, 0.473, 0.575)	0.607	(0.437, 0.668, 0.776)
	SPA-PV		SPA-IV		SPA-CV	
	Mean	Quartiles	Mean	Quartiles	Mean	Quartiles
\hat{H}		Omitted	0.116	(0.059, 0.118, 0.170)	0.130	(0.075, 0.124, 0.180)
$\hat{\chi}$			0.357	(0.000, 0.000, 0.822)	0.484	(0.081, 0.401, 0.954)
Estimates using only first 20 periods						
	FPA-PV		FPA-IV		FPA-CV	
	Mean	Quartiles	Mean	Quartiles	Mean	Quartiles
\hat{H}	0.324	(0.287, 0.341, 0.381)	0.236	(0.169, 0.229, 0.303)	0.258	(0.215, 0.271, 0.303)
$\hat{\chi}$	—	—	0.503	(0.000, 0.490, 1.000)	0.497	(0.010, 0.402, 0.895)
$\hat{\alpha}$	0.502	(0.280, 0.490, 0.615)	0.209	(0.090, 0.358, 0.605)	0.462	(0.338, 0.496, 0.719)
	SPA-PV		SPA-IV		SPA-CV	
	Mean	Quartiles	Mean	Quartiles	Mean	Quartiles
\hat{H}		Omitted	0.117	(0.057, 0.112, 0.162)	0.135	(0.088, 0.116, 0.173)
$\hat{\chi}$			0.428	(0.000, 0.163, 1.000)	0.600	(0.280, 0.653, 1.000)
Estimates using only last 20 periods						
	FPA-PV		FPA-IV		FPA-CV	
	Mean	Quartiles	Mean	Quartiles	Mean	Quartiles
\hat{H}	0.357	(0.318, 0.389, 0.408)	0.295	(0.228, 0.322, 0.363)	0.310	(0.276, 0.335, 0.352)
$\hat{\chi}$	—	—	0.571	(0.061, 0.635, 1.000)	0.513	(0.109, 0.524, 1.000)
$\hat{\alpha}$	0.580	(0.394, 0.507, 0.686)	0.483	(0.341, 0.465, 0.623)	0.730	(0.531, 0.843, 0.943)
	SPA-PV		SPA-IV		SPA-CV	
	Mean	Quartiles	Mean	Quartiles	Mean	Quartiles
\hat{H}		Omitted	0.129	(0.077, 0.143, 0.185)	0.139	(0.083, 0.128, 0.199)
$\hat{\chi}$			0.321	(0.000, 0.000, 0.827)	0.404	(0.000, 0.263, 0.922)

Table 6: Descriptive statistics of bid function parameter estimates, all treatments. Provided are the mean values of each parameter, and the quartiles of the distribution, where each observation is one participant.

	FPA			SPA		
	PV	IV	CV	PV	IV	CV
Signal	0.58** (0.160)	0.18** (0.018)	0.36** (0.108)	-0.11** (0.020)	-0.20 (0.242)	-0.11 (0.089)
1/ <i>t</i>	3.60 (2.208)	4.17** (1.491)	2.06 (1.789)	6.31** (1.007)	4.15** (1.171)	3.22** (1.039)
Female	-2.222 (1.190)	-1.92 (2.089)	-1.62 (1.765)	-0.94 (0.551)	0.07 (0.927)	0.42 (1.480)
English (lang)	-1.62 (1.284)	-6.06** (1.439)	-3.61 (1.887)	-0.69 (0.410)	-2.57 (2.778)	1.22 (1.708)
Economics	1.07 (1.315)	-5.40** (2.101)	-1.36 (1.586)	0.59 (1.442)	-1.75 (1.110)	0.63 (1.369)
Constant	7.37** (1.010)	14.81** (2.827)	11.89** (3.178)	6.97** (0.484)	12.35** (4.512)	7.76** (0.228)
Observations	1280	1280	1280	1280	1280	1280

Table 7: Determinants of time to submit bids. Dependent variable is response time in seconds. Female, English (lang) and Economics are dummy variables for women, native English speakers, and students majoring in economics; *t* indexes auction periods (1 to 30). Random-effects regression, with standard errors clustered by sessions; robust standard errors are reported in parentheses. ** indicates coefficient is significantly different from zero at the .01 level; * at .05.