Tu mihi soli places: An experiment on the competitiveness of all-pay auctions with private information[1]

Lucas Rentschler
Jon M. Huntsman School of Business
Utah State University
Logan, Utah
lucas.rentschler@usu.edu

Theodore L. Turocy
School of Economics and
Centre for Behavioural and Experimental Social Science
University of East Anglia
Norwich, United Kingdom
T.Turocy@uea.ac.uk

November 6, 2017

[1]**PRELIMINARY AND INCOMPLETE.** Please do not circulate or cite without permission. We thank conference participants at the 2015 SAET meetings in Cambridge, and seminar participants at Cardiff University, Durham University, Purdue University, University of East Anglia, and University of Exeter, for helpful comments. Turocy also thanks the College Board for setting the poem in question on the Advanced Placement examination circa 1990 and thereby inspiring this paper. All errors are the responsibility of the authors.

Abstract

In all-pay auctions with affiliated types, the implications of being of a "higher" type can be complex. Having a high assessment of the value of the prize is good news, but only if the other participants in the contest are not too likely also to have high assessments. If anticipated competition is strong, it is not clear whether a high or low bid will be a best response. In a laboratory experiment, we study behavior in both private-values and common-values settings with two contestants. We find general support for the comparative statics of Bayes-Nash equilibrium for private values. In contrast, behavior in common values settings in which bidders have very noisy information about the value of the prize differs greatly from the equilibrium predictions.

JEL Classifications: D44; D82; D72.

Keywords: contests, all-pay auctions, mixed strategies, affiliation, monotonicity.

1 Introduction

Atque utinam posses uni mihi bella videri!
Displiceas aliis: sic ego tutus ero.

And would that you could seem beautiful to me alone!
May you displease others: thus I will be safe. - Tibullus 3.19[1]

In this poem, Tibullus invokes an age-old trope of the romantic genre: The lovelorn youth, thinking of the object of his desires, wishes that other potential suitors would not find "his" girl attractive. The contest for her affections, then, would be less competitive, and he would not be faced with the decision whether to pursue her actively, or, as the poem continues, to risk envying the rival who might instead win her hand.[2]

Similar strategic reasoning applies in situations of economic interest. Consider an entrepreneur who is developing a product to enter a new market. Suppose that there will be network effects in this market so that the first entrant to the market is likely to obtain a dominant position and that higher research and development expenditures can hasten entry. If the entrepreneur observes information that suggests that the demand in the market is likely to be high, other things equal this makes competing to be the first entrant attractive. However, if the entrepreneur believes other potential entrants are likely to have made similar assessments, this information may be bad news, as it increases the risk of paying out those expenditures without winning the contest in the end.

These considerations distinguish contests such as the all-pay auction from winner-pay auction games. In a winner pay auction with affiliated types, having a "higher" type generally implies higher optimal bids, even if types are highly correlated, as the bid only needs to be paid conditional on winning. In the all-pay auction, the possibility of facing stiff competition makes low bids potentially attractive; low bids have a small probability of winning but cost little if they are unsuccessful.

In the all-pay auction with affiliated types drawn from a continuum, Krishna and Morgan (1997) provided a condition for the equilibrium to be monotonic in type. This condition, roughly speaking, requires that types are "not too affiliated;" that is, that the good news of having a higher type is not canceled out by the bad news that the contest is likely to be competitive. In the case of discrete types, Siegel (2014) proves a similar monotonicity result using the discrete analog to Krishna and Morgan's condition, but without other restrictions on the distribution of types. Siegel's

[1]The title is a play on an earlier line in this poem. The original line is *Tu mihi sola places*, "You alone please me." Our title means instead, "You are pleasing to me alone."

[2]While modern sensibilities may find this characterization of courtship to be at best inappropriate, there are many historical periods and cultures in which male suitors have framed the pursuit of a mate in such terms.

result formalizes the intuition that the conditions for monotonicity in all-pay auctions are distinct from those in winner-pay auctions (e.g., affiliation as in Milgrom and Weber 1982).

In all-pay auctions where the equilibrium is not monotonic, bidders face, roughly, the choice between "going for it" with a relatively high bid, or playing safe with a low one. Rentschler and Turocy (2016) provide a complete characterization of equilibria in the two-bidder case, including examples of equilibria with complex structures involving randomization over disjoint intervals of bids. The construction in Rentschler and Turocy also demonstrates that the complexity of the process of computing an equilibrium can be much higher in settings where the equilibrium is not monotonic.

All-pay auctions with discrete types generally have equilibria in mixed strategies. By their nature mixed-strategy equilibria depend sensitively on the assumption that all players use best responses. This presents a challenge to the use of these equilibria as a behavioral prediction. This challenge is intensified in the case of all-pay auction games without monotonic equilibria, as the complexity of the equilibria themselves, and of the computational burden in identifying them, undermines their plausibility as an exact description of behavior.

We design an experiment which asks whether the qualitative prediction embodied by the Krishna-Morgan-Siegel monotonicity condition nevertheless holds in the presence of non-equilibrium play. Is it definitely "good news" to have a "high" type when types are not strongly correlated, but less so (or not at all) when correlation is high? Further, as the monotonicity condition applies equally to private-values and common-values contests, does the structure of how types related to valuations make a difference? We find that for private values, it is indeed the case that "high" types do better when the values are not very correlated, but not when they are highly correlated. However, for common values, the equilibrium prediction for the effect of accuracy of information is not found in our data.

The current experimental literature on all-pay auctions with incomplete information is quite small, and predominantly focuses on the case of independent types. One exception is Grosskopf et al. (2014), which considers common-value all pay auctions with conditionally independent signals. The current paper differs in that we vary the degree of correlation, and directly compare behavior to equilibrium predictions.[3] Noussair and Silver (2006) examine an independent private values environment with six contestants. They find that overbidding is high, such that many participants were bankrupt by the end of the session. Hörisch and Kirchkamp (2010) examine both wars of attrition and all-pay auctions with independent private marginal costs of expenditure, and also find significant overbidding in all-pay auctions. Müller and Schotter (2010) examine all-pay auctions with four contestants and either linear or quadratic bidding costs. A multiplicative cost

[3] Athey (2001) shows that in the symmetric environment studied in Grosskopf et al. (2014) the single crossing property fails.

parameter is independently drawn and is private information. In both cost environments, expenditures are bimodal, and, on average, above Nash predictions. Hyndman et al. (2012) vary the feedback provided to participants to determine whether anticipated regret will affect bidding behavior. In an independent private values framework, they find that, regardless of whether on not the winning expenditure is revealed, bids are, on average, above Nash predictions. However, this tendency is even more pronounced when the winning bid is revealed, suggesting that anticipated regret may have an effect on behavior.

The remainder of the paper is organized as follows. Section 2 describes the theoretical framework we study. Section 3 details our experimental design. Section 4 presents our results, and Section 5 contains a discussion of these results.

2 Theory

The general framework follows that used in Krishna and Morgan (1997), Siegel (2014), and Rentschler and Turocy (2016). There are two contestants who compete for a single, indivisible prize. Each contestant has a type drawn from a set T, which is their private information; in our experiment, we specialize to the case of two types, with $T = \{t_L, t_H\}$. After observing their types, contestants simultaneously submit non-negative bids. Each bid is irrevocably sunk, but only the bidder with the higher of the bids wins the prize. In the event of a tie, the winner is determined by a fair randomization.

A bidder's valuation of the prize may depend on both her and her opponent's types. If a bidder's type is $t_k \in T$ and the opponent's type is $t_l \in T$, the posterior expected value of the prize is $V_{k,l}$. Given that a bidder is of some type $t_k \in T$, the conditional probability that the other bidder is type $t_l \in T$ is $h_{l|k}$.

We consider symmetric Bayes-Nash equilibria. Rentschler and Turocy (2016) show that all symmetric equilibrium will involve piecewise uniform randomization for both types, and that there will be no atoms in the distribution of bids for either type. We denote such a behavior strategy as π, which assigns to each type $t_k \in T$ a corresponding probability density function π_k over bids. The expected monetary payoff to a contestant of type t_k who bids b against an opponent who plays according to behavior strategy π is

$$u_k(b|\pi) = \sum_{l:t_l \in T} h_{l|k} V_{k,l} \left[\int_0^b \pi_l(b) \right] - b$$

where ties are neglected because they occur with zero probability. For notational compactness we write $\psi_{l|k} \equiv h_{l|k} V_{k,l}$, which defines a matrix ψ with rows indexed by t_l and columns by t_k.

Siegel (2014) showed that the qualitative properties of the Bayes-Nash equilibria of this game depend on the structure of the matrix ψ. Specifically, if the entries of each row are strictly monotonic, then there is a unique Bayes-Nash equilibrium which is stochastically monotonic. In our setting, stochastic monotonicity means that all bids submitted by type t_H are greater than those submitted by type t_L.

The monotonicity condition on ψ ensures that being of type t_H is unambiguously "good news" for the bidder, relative to being of type t_L. The dependence of ψ on both the posterior expected value and the conditional probabilities of the opponent's type distinguishes this condition from the case of the (winner-pay) first-price auction, in which affiliation of values is a sufficient condition for monotonicity of the equilibrium (Milgrom and Weber, 1982; Wang, 1991). Krishna and Morgan (1997) note that the monotonicity condition rules out the case of values which are "too affiliated."

Our experimental design contrasts, in one dimension, the case in which receiving the higher type is unambiguously good news with the case in which the monotonicity condition is violated by values which are strongly positively correlated. In our environment, there are two possible states of the world, drawn from $\Omega = \{\omega_L, \omega_H\}$ with equal probability. The state of the world is not revealed directly to the contestants. Instead, conditional on the state being ω_L, each contestant's type is t_L with probability p, and t_H with probability $1 - p$; similarly, if the state is ω_H, each contestant's type is t_H with probability p and t_L with probability $1 - p$. The realization of contestant types is done independently, conditional on the underlying state.

The monotonicity condition does not directly restrict how types map into valuations. In a second dimension our design contrasts the case of pure private values with pure common values. With private values, the valuation of a contestant is equal to her type. With common values, the valuation of a contestant is equal to the state of the world. Importantly, the qualitative implication of the monotonicity condition is the same for both private and common value settings. If p is sufficiently close to $\frac{1}{2}$, the monotonicity condition is satisfied, and the resulting equilibrium is stochastically monotonic. If p is sufficiently close to 1, the monotonicity condition is violated, and in equilibrium a contestant of type t_H will lose to one of type t_L with positive probability.

In this setting, the matrix ψ is

$$\psi = \begin{bmatrix} \left[p^2 + (1-p)^2\right] V_{L,L} & 2p(1-p)V_{H,L} \\ 2p(1-p)V_{L,H} & \left[p^2 + (1-p)^2\right] V_{H,H} \end{bmatrix}$$

Monotonicity of the second row of the matrix follows if $V_{H,H} > V_{L,H}$; therefore, in order for the equilibrium to be monotonic, the parameters must satisfy

$$\frac{V_{L,L}}{V_{H,L}} < \frac{2p(1-p)}{1 - 2p(1-p)} \tag{1}$$

For private values, $V_{L,L} = \omega_L$ and $V_{H,L} = \omega_H$, and so

$$\frac{\omega_L}{\omega_H} < \frac{2p(1-p)}{1-2p(1-p)} \tag{2}$$

For common values, the conditional posterior values are given by

$$\frac{V_{L,L}}{V_{H,L}} = \frac{2}{p^2+(1-p)^2} \times \frac{p^2\omega_L + (1-p)^2\omega_H}{\omega_L+\omega_H}. \tag{3}$$

Combining (1) and (3) and simplifying gives

$$\begin{aligned}\frac{2}{p^2+(1-p)^2} \times \frac{p^2\omega_L + (1-p)^2\omega_H}{\omega_L+\omega_H} &< \frac{2p(1-p)}{1-2p(1-p)} \\ \frac{\omega_L}{\omega_H} &< \frac{1-p}{p}.\end{aligned} \tag{4}$$

To jointly satisfy (2) and (4), in our experiment, we choose $\omega_L = 15$ and $\omega_H = 30$. For private values, (2) requires $p < \frac{1}{2} + \sqrt{\frac{1}{12}} \approx 0.789$ for the equilibrium to be monotonic, and for common values (4) requires $p < \frac{2}{3}$.

We choose $p = 0.6$ for low correlation treatments, in which the equilibrium prediction is stochastically monotonic, and $p = 0.9$ for high correlation treatments. In what follows we refer to private value environments with low correlation between types ($p = 0.6$) as the LPV case. Similarly, we refer to private values with a high degree of correlation ($p = 0.9$) as the HPV case. When the valuation structure is common-value we use LCV and HCV for the case of $p = 0.6$ and $p = 0.9$, respectively.

Calculation of the equilibria for each condition is straightforward using the results in Rentschler and Turocy (2016). For the LPV case, the equilibrium behavior strategy is given by

$$\pi_L^{LPV}(b) = \begin{cases} \frac{5}{39} & \text{if } b \in \left[0, 7\frac{4}{5}\right] \\ 0 & \text{otherwise} \end{cases}$$

$$\pi_H^{LPV}(b) = \begin{cases} \frac{5}{78} & \text{if } b \in \left[7\frac{4}{5}, 23\frac{2}{5}\right] \\ 0 & \text{otherwise.} \end{cases}$$

Note that, since the monotonicity condition is satisfied, the supports of the two equilibrium densities do not overlap. As such, the high type will win with probability one when pitted against an opponent of the low type. The high type earns a positive expected payoff, while the low type's expected payoff is zero.

Turning attention to the HPV case, the equilibrium behavior strategy is

$$\pi_L^{HPV}(b) = \begin{cases} \frac{73}{960} & \text{if } b \in \left[0, 13\frac{11}{73}\right] \\ 0 & \text{otherwise} \end{cases}$$

$$\pi_H^{HPV}(b) = \begin{cases} \frac{23}{960} & \text{if } b \in \left[0, 13\frac{11}{73}\right] \\ \frac{5}{123} & \text{if } b \in \left[13\frac{11}{73}, 30\right] \\ 0 & \text{otherwise.} \end{cases}$$

Note that the support of the high type's density now overlaps with that of the low type. As a result, the high type will now lose to the low type with positive probability. Also, since zero is in the support of π_H^{HPV}, the corresponding equilibrium expected payoff is also zero. Additionally, note that relative to the LPV case, the upper bound of the support of the high type's density is higher. Thus, the high type is indifferent between bidding very aggressively ("going for it"), and bidding very cautiously.

In addition to the LPV and HPC cases, we also consider the case of independent private values (IPV), in which $p = 0.5$. Under independent private values the equilibrium behavior strategy is

$$\pi_L^{IPV}(b) = \begin{cases} \frac{2}{15} & \text{if } b \in \left[0, 7\frac{1}{2}\right] \\ 0 & \text{otherwise} \end{cases}$$

$$\pi_H^{IPV}(b) = \begin{cases} \frac{1}{15} & \text{if } b \in \left[7\frac{1}{2}, 22\frac{1}{2}\right] \\ 0 & \text{otherwise.} \end{cases}$$

Note that equilibrium behavior in the IPV case is qualitatively similar to that of the LPV case. We opted to study the IPV case because we were concerned that conveying the concept of correlated types to our experimental subjects would be a challenge. By including the IPV case, we have a benchmark case with uncorrelated types we can compare the LPV case to. If behavior does not dramatically deviate between these two cases, despite one having (weakly) positively correlated types, we can be guardedly optimistic that subjects understood how we introduced the positive correlation. Another reason to include the IPV case is that this parameterization allows us to benchmark our results against the results of Noussair and Silver (2006).

Turing attention to environments with common values, the equilibrium behavior strategy for

	PV			CV	
	IPV	LPV	HPV	LCV	HCV
	$p = 0.5$	$p = 0.6$	$p = 0.9$	$p = 0.6$	$p = 0.9$
Mean bid of t_L	3.75	3.90	6.58	5.10	7.06
SD bid of t_L	2.17	2.25	3.80	2.94	4.08
Mean bid of t_H	15.00	15.60	16.85	16.80	15.44
SD bid of t_H	4.33	4.50	4.53	3.78	4.12
Mean bid of t_H - Mean bid of t_L	11.25	11.70	10.27	11.70	8.38
Higher type wins	100%	100%	84.2%	100%	79.4%

Table 1: Summary of equilibrium predictions.

the LCV case is

$$\pi_L^{LCV}(b) = \begin{cases} \frac{5}{51} & \text{if } b \in \left[0, 10\frac{1}{5}\right] \\ 0 & \text{otherwise} \end{cases}$$

$$\pi_H^{LCV}(b) = \begin{cases} \frac{5}{66} & \text{if } b \in \left[10\frac{1}{5}, 23\frac{2}{5}\right] \\ 0 & \text{otherwise.} \end{cases}$$

Predicted behavior is qualitatively similar to that of LPV. Equilibrium behavior is stochasticly monotonic in type, the low type has an expected profit of zero, while that of the high type is positive. However, notice that the upper bound of the support of π_L^{LCV} is higher than that of π_L^{LPV}. This is driven the uncertainty regarding the value of winning in the LCV case.

For the high correlation common-value case the equilibrium behavior strategy is

$$\pi_L^{HCV}(b) = \begin{cases} \frac{17}{240} & \text{if } b \in \left[0, 14\frac{2}{17}\right] \\ 0 & \text{otherwise} \end{cases}$$

$$\pi_H^{HCV}(b) = \begin{cases} \frac{7}{240} & \text{if } b \in \left[0, 14\frac{2}{17}\right] \\ \frac{20}{489} & \text{if } b \in \left[14\frac{2}{17}, 28\frac{1}{2}\right] \\ 0 & \text{otherwise.} \end{cases}$$

Behavior in the HCV case is not predicted to substantially differ from t hat of the HPV case. Both types have an expected payoff of zero; the high type is not better off than the low type. Further, equilibrium behavior is not stochastically increasing in type.

Table 1 summarizes the key quantities predicted by the Bayes-Nash equilibrium for all five parameterizations, which we will use as the basis for our subsequent data analysis.

3 Experimental design

We report on a total of twenty-five experimental sessions conducted at Universidad Francisco Marroquín. Each session used eight participants drawn from the subject pool maintained by Centro Vernon Smith de Economía Experimental. None of the subjects participated in more than one session, and none had prior experience in contest games. All interaction was mediated via computer, and no participant IDs or other identifying information was available regarding the other participants in a session.[4]

We vary treatments between sessions. There are a total of five treatments, which vary whether the prize had a private or common value, and the correlation between types. We conducted five sessions for each treatment. Within a session, participants were paired at random at the start of the session, and remained in the same pair throughout the session. Therefore, we have twenty independent observations at the session-pair level.

There were forty periods in a session. In each period, the contestants participated in an all-pay sealed-bid auction. Each received a type drawn from the set $\{15.00, 30.00\}$.[5] We used the same realizations of types and states of the world for each session within a treatment.

At the beginning of each auction period, participants' computer screens displayed their type and a five second countdown clock. After the countdown participants could select a bid using a slider device.[6] Bids could be submitted in units of 0.30.[7] Bids were specified by clicking a point along the slider, and then confirming with a button click. The period lasted at least forty seconds, or until five seconds after the last contestant had submitted her bid. The results of the auctions were displayed for fifteen seconds, after which the next period began.[8]

We extended the feedback mechanism employed for winner-pay auctions by Turocy and Cason (2015) to the case of all-pay auctions. In addition to providing feedback in the same graphical frame as used by the participant to choose a bid, the feedback screen also provided information on the results of the auction from the perspective of the other participant. Figure 1 shows sample

[4]The full text of the instructions, as well as translations from the original Spanish, is available as a separate appendix.

[5]All types, bids, and earnings were expressed directly in Guatemalan Quetzales; we avoided the use of in-lab currency units and associated exchange rates.

[6]This device was used previously in winner-pay auctions in Turocy et al. (2007), Turocy and Watson (2012), and Turocy and Cason (2015).

[7]We selected the granularity of the type and bid spaces to be fine enough that ties in bids would be relatively unlikely, making the continuous approximation reasonable. The equilibrium in the discrete game exhibits a "sawtooth" alternation in the probabilities assigned to alternate bid levels, similar to that shown in discrete Blotto games by Hart (2008). We neglect this alternation and use the continuous equilibrium as our baseline, but our results would be unchanged by using the discretized equilibrium.

[8]We designed this pacing of the auction periods to control for the participants' opportunity cost of time, as no participant could make the session conclude faster by making choices faster, as well as to make the private type and the feedback process salient. Few contestants took more than forty seconds to submit a bid in any period.

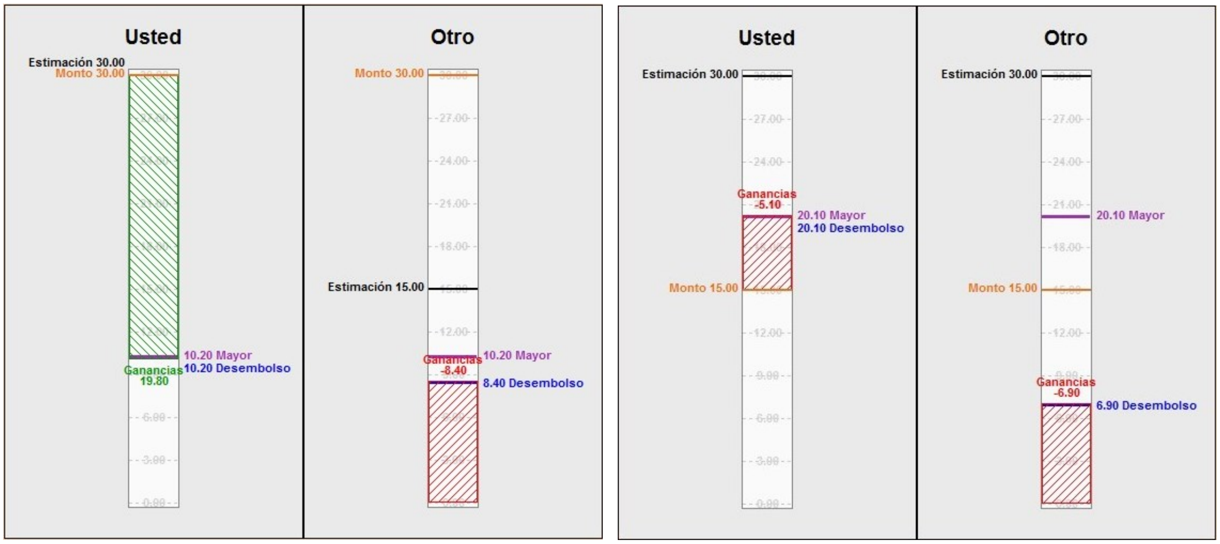

(a) Positive earnings

(b) Negative earnings

Figure 1: Structure of feedback for a typical period.

screenshots with typical outcomes in the cases of positive and negative profit by the contestant after winning the auction.

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

Participants received an initial balance of $Q120.00 \approx US\$15.67$. At the end of the session, ten out of the forty periods were selected at random for payment. Therefore, a string of negative earnings in early periods would not necessarily cause bankruptcy, eliminating a design complication which can arise both in winner-pay common-value experiments as well as contests more generally. The initial balances were set large enough that the chances of negative total earnings were small, even under assumptions of very aggressive expenditure choices. In the event that the ten selected periods resulted in a negative balance, participants received no money for the session.[9] On average, participants earned $Q100.83$, for sessions which lasted just over an hour, inclusive of instructions.

4 Results

We begin with a comparison of aggregate distributions of bids against the benchmark equilibrium prediction of piecewise-uniform randomization. Figure 2 presents cumulative empirical distribu-

[9]Only 8 of the 200 participants left with zero earnings: three in LCV, three in HCV, and one each in LPV and HPV.

Quantity	Statistic	PV			CV	
		IPV $p = 0.5$	LPV $p = 0.6$	HPV $p = 0.9$	LCV $p = 0.6$	HCV $p = 0.9$
Mean bid of t_L	Median over pairs	6.48	7.23	10.16	11.13	10.10
	Pairs > NE	17**	19***	16**	18***	15*
SD bid of t_L	Median over pairs	4.94	5.20	4.88	6.61	5.20
	Pairs > NE	19***	18***	15*	20***	14
Mean bid of t_H	Median over pairs	19.23	19.47	20.49	16.54	19.24
	Pairs > NE	19***	17**	16*	7	12
SD bid of t_H	Median over pairs	5.56	6.08	7.65	6.37	7.91
	Pairs > NE	16*	17**	18***	19***	19***
(Mean bid of t_H -	Median over pairs	12.80	12.14	11.48	4.47	6.92
Mean bid of t_L)	Pairs > NE	11	11	11	0***	7
Higher type wins	Mean over pairs	91.3%	86.2%	79.5%	70.8%	78.6%
Sum of bids	Median over pairs	26.54	26.73	32.31	26.76	30.39
	Pairs > NE	19***	18***	16*	13	15*
Earnings of t_L	Median over pairs	-2.46	-2.38	-3.43	-3.45	-3.22
	Pairs > 0	3**	3**	5*	5*	6
Earnings of t_H	Median over pairs	2.06	0.53	-4.60	-1.07	-4.56
	Pairs > 0	16*	12	4**	7	6

Table 2: Comparison of observed quantities against equilibrium predictions. The unit of independent observation is the pair; there are 20 independent pairs per treatment. For the test of whether the median value is different from the Nash equilibrium prediction, * denotes significant at 5%, ** at 1%, and *** at 0.1%.

tions of bids for private values treatments, and 3 the same for common-values. For all treatments, we present the distributions separately for the first 20 periods (first half) and last 20 periods (second half) of the experiment. The benchmark equilibrium distributions are provided for reference.

Some general patterns emerge from the overall distributions. Bidding is generally more aggressive than the equilibrium prediction in all treatments, throughout the course of the experiments. There is some evidence of adjustment, in that bids distributions in the second half are generally closer to the equilibrium prediction than in the first half, with convergence in IPV and HPV for t_L being visually notable. Bids in private values treatments show substantial responsiveness to type, with bid distributions for t_H being well to the right of t_L. Bids in common values treatments are less responsive to type, especially in the LCV treatment, which contrasts with the equilibrium prediction of bidding monotonic in type.

We turn to a formal quantitative analysis. Each pair of bidders remained matched throughout the experiment. We therefore use the pair as the unit of independent observation, and conduct statistical testing accordingly. We present a summary of each participant's bid distributions in

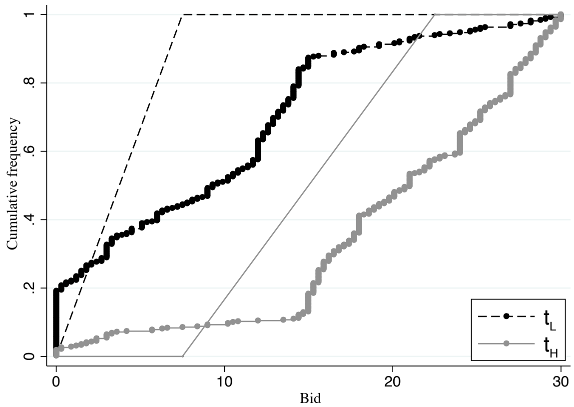

(a) IPV, first 20 periods

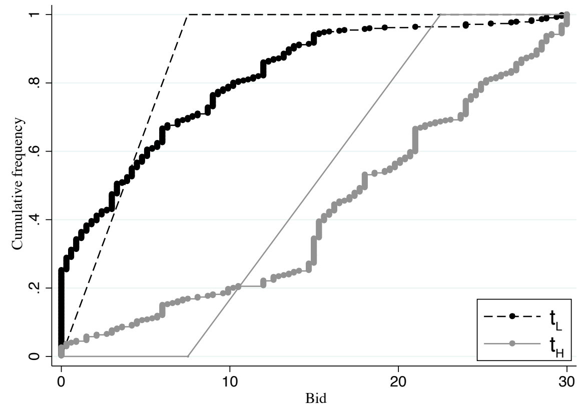

(b) IPV, last 20 periods

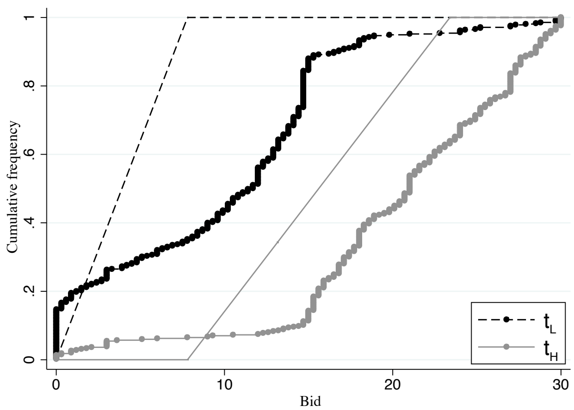

(c) LPV, first 20 periods

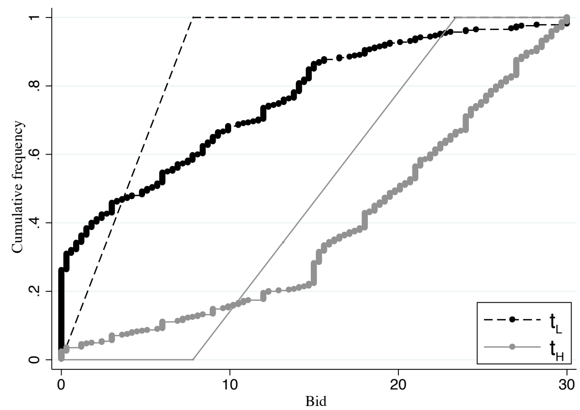

(d) LPV, last 20 periods

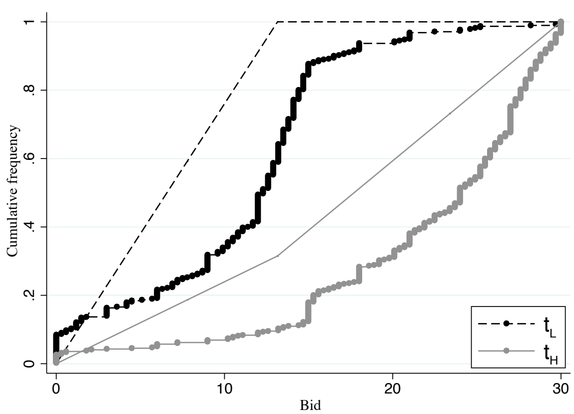

(e) HPV, first 20 periods

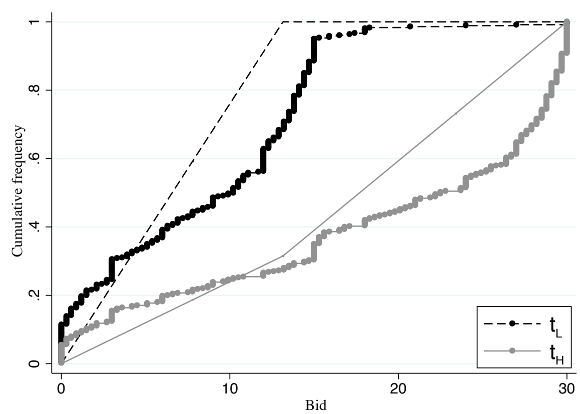

(f) HPV, last 20 periods

Figure 2: CDF of expenditures, first 20 periods versus last 20 periods, private values

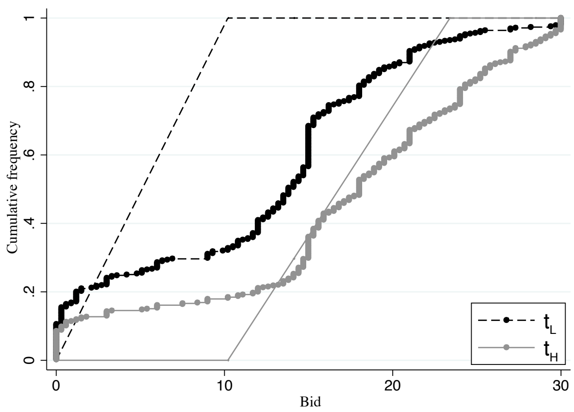

(a) LCV, first 20 periods

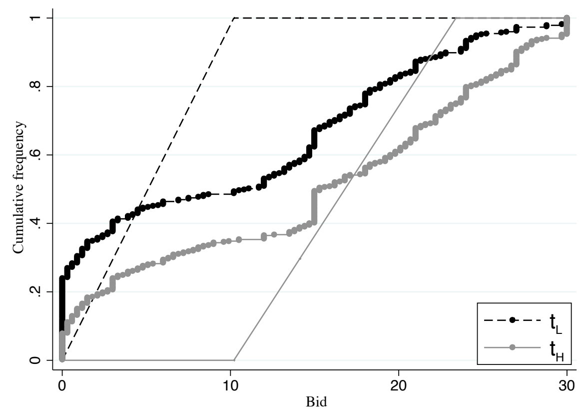

(b) LCV, last 20 periods

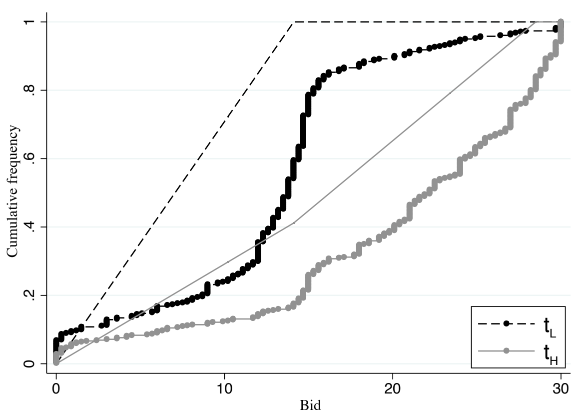

(c) HCV, first 20 periods

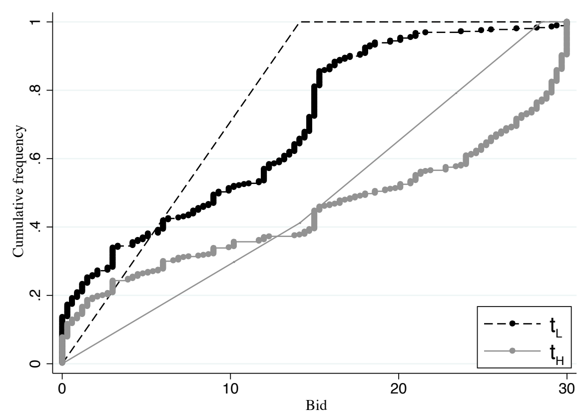

(d) HCV, last 20 periods

Figure 3: CDF of expenditures, first 20 periods versus last 20 periods, common values

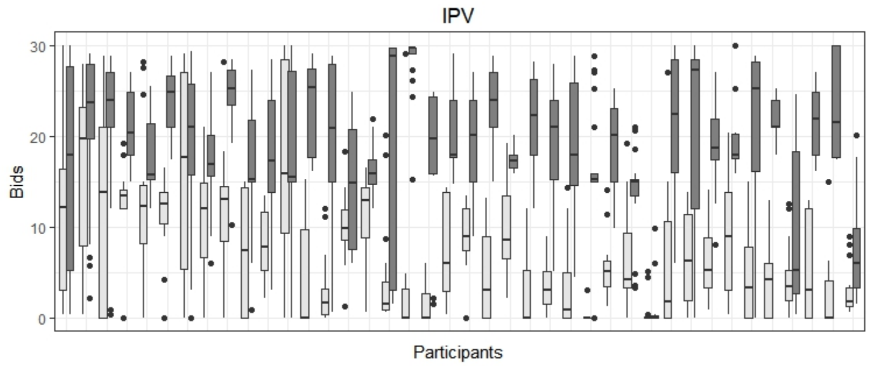

(a) IPV

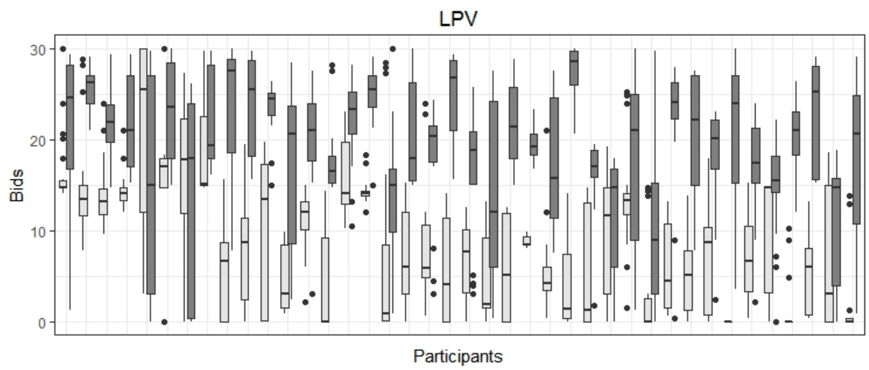

(b) LPV

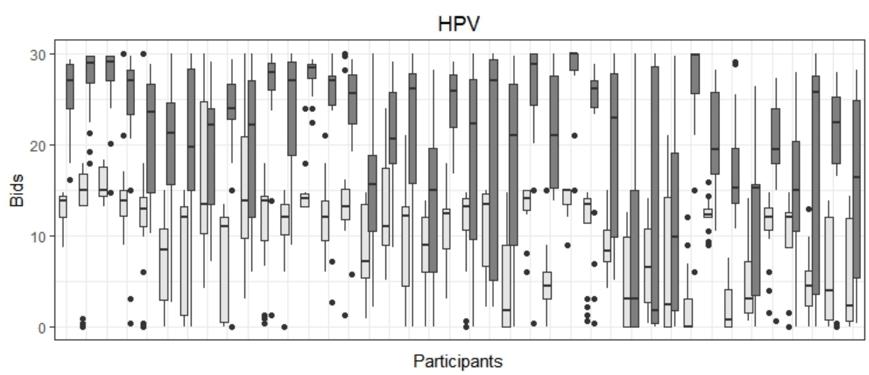

(c) HPV

Figure 4: Bidder-by-bidder boxplots of distribution of bids, private values treatments. The behavior of each bidder is summarized by two adjacent boxplots, corresponding to her behavior when she was the low type (light grey) and when she was the high type (dark grey). Bidders are sorted from left to right by average payoffs across the 40 periods of the experiment, with relatively low payoffs on the left.

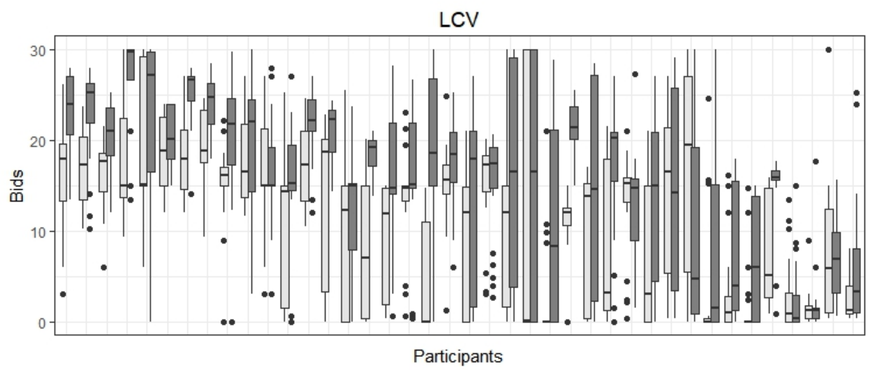

(a) LCV

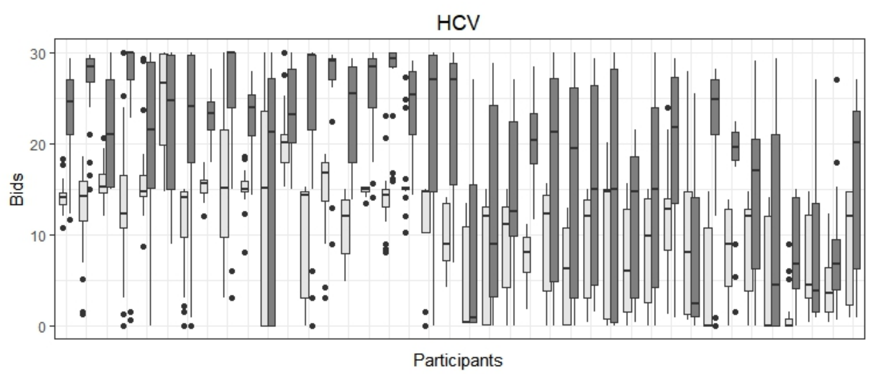

(b) HCV

Figure 5: Bidder-by-bidder boxplots of distribution of bids, private values treatments. The behavior of each bidder is summarized by two adjacent boxplots, corresponding to her behavior when she was the low type (light grey) and when she was the high type (dark grey). Bidders are sorted from left to right by average payoffs across the 40 periods of the experiment, with relatively low payoffs on the left.

Quantity	Statistic	PV			CV	
		IPV $p=0.5$	LPV $p=0.6$	HPV $p=0.9$	LCV $p=0.6$	HCV $p=0.9$
Mean Δbid of t_L	Median over pairs	-1.89	-2.78	-2.48	-0.26	-4.43
	Pairs < 0	19***	14	17**	11	16*
Mean Δbid of t_H	Median over pairs	-3.00	-1.19	-2.43	-0.60	-4.42
	Pairs < 0	16*	13	14	11	13

Table 3: Comparison of mean bids by type in second half versus first half of experiment. The unit of independent observation is the pair; there are 20 independent pairs per treatment. For the test of whether the median value is different from the Nash equilibrium prediction, * denotes significant at 5%, ** at 1%, and *** at 0.1%.

Figure 4 for private values and Figure 5 for common values. There is a boxplot for the distribution of bids with the low type $t_L = 15$ and the high type $t_H = 30$ for each participant.

Table 2 summarizes the observations on the observed measures of behavior for which the equilibrium benchmarks appear in Table 1.

Result 1. *Bids generally exceed the Nash equilibrium prediction across treatments, in accordance with most results previously reported in experimental contests. Bidders vary their bids more than the equilibrium randomization predicts.*

Support. Table 2 reports summaries bids as a function of type. For each pair, we measure the bids by computing the mean bid submitted by the participants in the pair when they received the low type, and when they received the high type. In the table we report the median of these measures across all pairs, and the number of pairs for which the measure exceeds the Nash equilibrium prediction. We test the equilibrium prediction by testing the proportion of pairs whose median bid exceeds the equilibrium prediction. In PV, we reject at 5% the null hypothesis of the equilibrium prediction for the bids both types in all correlation conditions. In CV, mean bids are greater than the equilibrium prediction for t_L, but not for t_H.

Turning to within-participant variation in bids, for each pair we measure the standard deviation of bids conditional on type by computing the standard deviation of bids conditional on t_k for each bidder in the pair, and then taking the average of those measures. In the table we report the median of these measures across all pairs, and the number of pairs for which the measure exceeds the Nash equilibrium prediction. The standard deviation of bids exceeds the Nash prediction for both types across all treatments but t_L in HCV. We therefore conclude that the aggressive average bidding does not arise solely from a simple upwards shift of bid distributions. □

Result 2. *Bids are typically lower during the second half of the experiment.*

Quantity	Statistic	PV			CV	
		IPV $p = 0.5$	LPV $p = 0.6$	HPV $p = 0.9$	LCV $p = 0.6$	HCV $p = 0.9$
Within-bidder monotonicity	Mean over pairs	87.7%	87.0%	84.5%	69.3%	77.4%
	SD over pairs	(2.4%)	(2.8%)	(2.1%)	(2.1%)	(2.9%)
	Equilibrium	100.0%	100.0%	84.2%	100.0%	79.4%

Table 4: Within-subjects measure of overlap between high-type and low-type bid distributions, by treatment. Twenty independent observations per treatment.

Support. Table 3 presents measures of adjustment of bids over the course of the experiment. For each type $t_k \in \{t_L, t_H\}$, we compute for each pair the mean bid in the last 20 periods when the participants received type t_k, and subtracted the mean bid in the first 20 periods with type t_k. The table reports the median across pairs of these measures, and counts of the number of pairs for which bids in the second half were lower. For all treatments, we observe a reduction of bids in the second half in a majority of pairs. The effect is only strong enough to be statistically significant for both types in IPV, and for t_L in both high-correlation conditions HPV and HCV. □

Result 3. *Stochastic monotonicity of expenditures is rejected in both common and private values. A fortiori, in common values, expenditures are actually more dependent on type when types are highly correlated, which is the opposite of the comparative statics prediction of equilibrium.*

Support. The most direct test of the prediction of monotonicity is to measure the proportion of cases in which a bidder with type t_L outbids a bidder with type t_H. Table 2 provides this measure, where the proportion is computed for each independent pair, and the average of those pair-level proportions is reported. The proportion of cases in which t_H wins is less than the equilibrium prediction in all treatments, although this observation is not on its own particularly surprising in IPV, LPV, and LCV due to the strong equilibrium point prediction of 100%. More instructive is the comparative static prediction: t_H is predicted to win less frequently against t_L in HPV relative to IPV or LPV, and in HCV relative to LCV. With private values, we indeed find that t_H wins less frequently against t_L in HPV than in IPV or LPV, although we find a substantial difference between IPV and LPV. However, with common values, the comparison goes in the direction opposite the equilibrium prediction: t_H wins against t_L more often in HCV than LCV.

The empirical winning percentage measure may be picking up heterogeneity across participants as much as non-monotonicity; some cases in which a type t_L outbids type t_H could be because the two participants are using different strategies, even though both participants' strategies might be monotonic. We therefore develop a complementary within-participant measure of monotonicity. For each participant, let B_H be their sequence of bids in periods in which they were type t_H, and

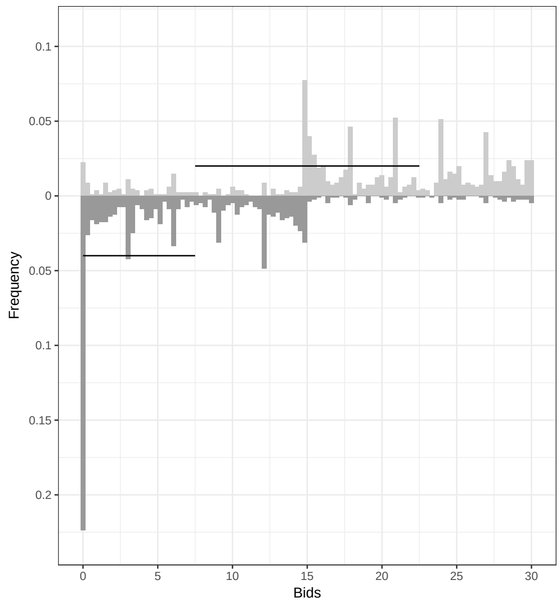

(a) IPV

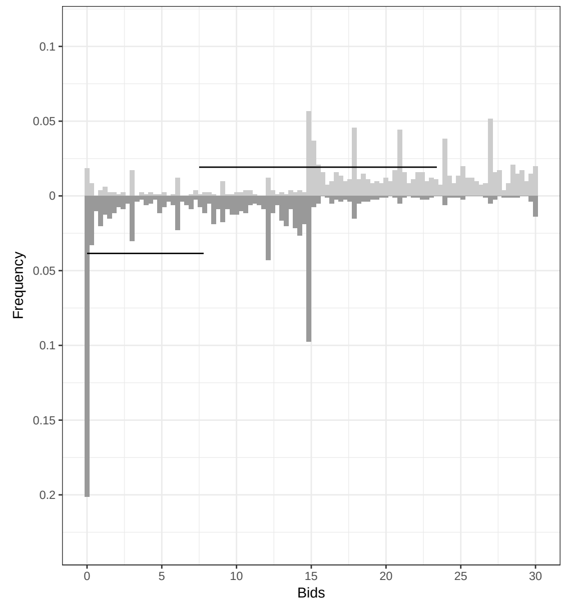

(b) LPV

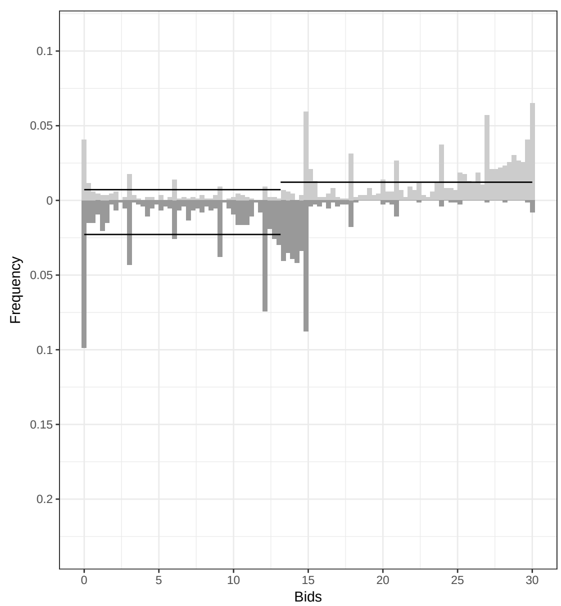

(c) HPV

Figure 6: Stacked histograms of bids by type in private value environments. Bids by high types (light grey) are on top, while bids by low types are on the bottom. The equilibrium bid densities are represented by solid black lines.

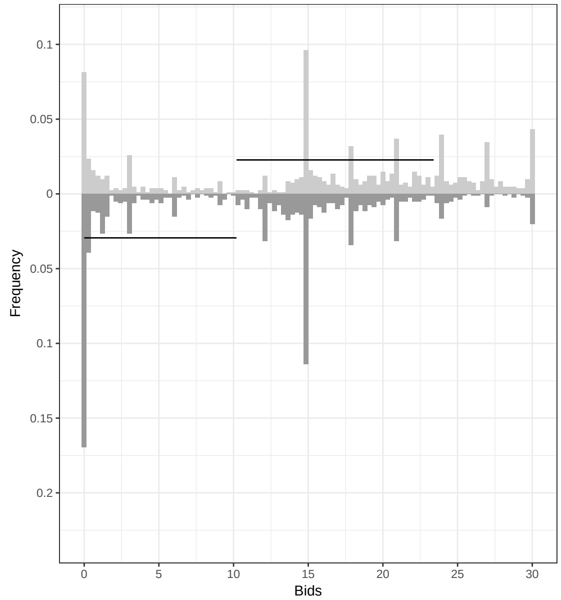

(a) LCV

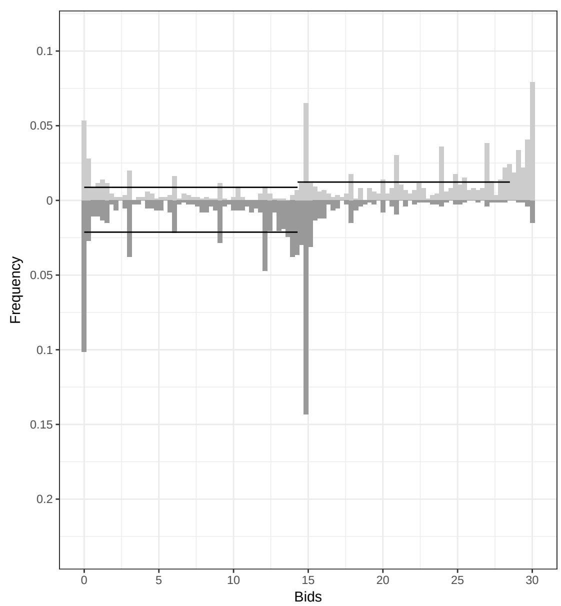

(b) HCV

Figure 7: Stacked histograms of bids by type in common value environments. Bids by high types (light grey) are on top, while bids by low types are on the bottom. The equilibrium bid densities are represented by solid black lines.

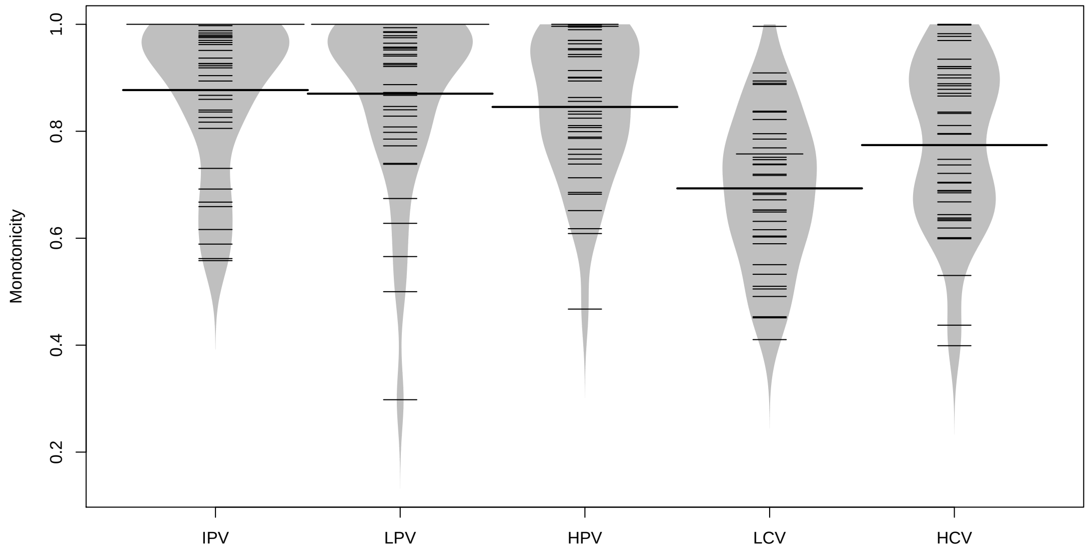

Figure 8: Beanplots of within-participant monotonicity by treatment.

B_L their sequence of bids in periods in which they were type t_L. We then pair up each entry in B_H with each entry in B_L, and count how frequently the bid from B_H is the higher. If a participant were following any stochastically monotonic strategy, this measure would be 100%. Lower values provide a measure of the extent of overlap in the participant's strategy. For each pair, we take the average of the overlap measures as the measure of monotonicity within the pair. Table 4 presents the average of this measure across all pairs, as well as the equilibrium predictions. The relative performance of treatments by this measure is identical to that obtained by directly computing empirical winning percentages. □

The monotonicity condition on ψ encodes the notation that being type t_H is unambiguously better news than t_L. In equilibrium in IPV, LPV, and LCV, t_H is predicted to receive positive payoffs on average. Nevertheless, the "good news" content of t_H can be evaluated without expecting behavior to be in equilibrium. We examine this first by considering the payoffs contingent on each type.

Result 4. *Earnings for type t_L are negative for all treatments. Earnings for t_H are positive only with private values and lower correlation of types (IPV and LPV).*

Support. Table 2 includes per-period net earnings for each type for all treatments. These are computed by taking, for each pair, the average earnings of the participants in the pair in periods in which they received t_L or t_H, respectively. For t_L, we can reject the null hypothesis of zero

Quantity	Statistic	PV			CV	
		IPV $p = 0.5$	LPV $p = 0.6$	HPV $p = 0.9$	LCV $p = 0.6$	HCV $p = 0.9$
$\pi^\star(t_H) - \pi^\star(t_L)$	Mean over pairs	8.53	6.02	1.27	1.28	1.85
	SD over pairs	(2.96)	(2.68)	(1.64)	(1.02)	(1.84)
	Median over pairs	9.21	6.39	0.50	1.44	1.23

Table 5: Comparison of best-response payoffs conditional on high versus low types, by treatment.

net earnings for all treatments except HCV. For t_H, in private values, the predicted comparative statics are found in the point predictions: participants have net positive earnings in IPV and LPV, but negative in HPV. Earnings for t_H are also higher in LCV than in HCV. Of particular note is that the most negative earnings are observed for t_H in HPV. □

Result 5. *The contests are more competitive than predicted by equilibrium, as measured by total bids, for all treatments.*

Support. Table 2 presents statistics on the sum of bids, measured again at the level of the pair. In all cases this measure exceeds the equilibrium prediction. Consistent with equilibrium predictions, the case of HPV is the most competitive, with the largest sum of bids by a large margin. □

We turn to a complementary approach for summarizing bidding behavior. The monotonicity condition on ψ is meant to capture the idea that it is better news about one's prospects in the contest if one receives type t_H than if one receives t_L. Result 4 measures this by the actual earnings participants received. However, participants are not in general best-responding to the behavior of the other participant in their pair.

The results so far are based on direct descriptions of the distributions of bids submitted by each participant. Another way to summarize behavior in a way that is strategically relevant is to ask what is the best response to each participant's empirical distribution of bids. For each participant i, we take as given the distribution of realized bids, as a function of their type, of the other participant in their pair over the course of the experiment. Given those distributions, we determine the best response bids conditional on type, $(b_i^\star(t_H), b_i^\star(t_L))$. Figure 10 shows these best response bids as a scatterplot, with each participant plotted at coordinates $(b_i^\star(t_H), b_i^\star(t_L))$. For all treatments, the best responses $b_\cdot^\star(t_L)$ cluster at or just above zero. This reflects the frequency with which bids of zero are submitted, as can also be seen in Figures 2 and 3. This bimodal pattern of bids, with frequent bids of zero combined with frequent aggressive bids, is also noted by Ernst and Thöni (2013) and Sheremeta (2013).

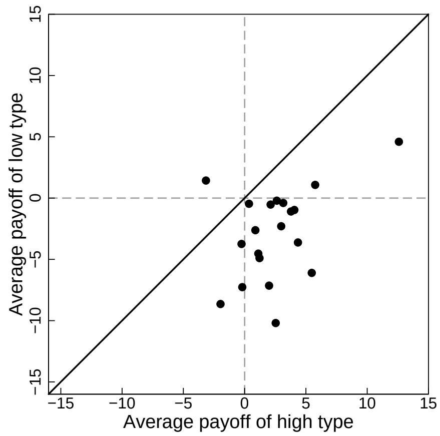

(a) IPV

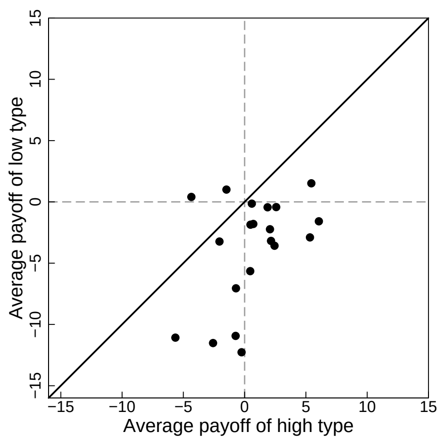

(b) LPV

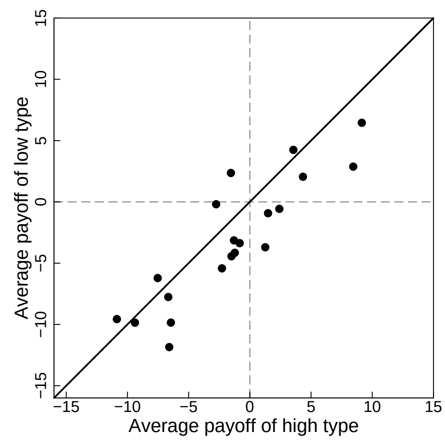

(c) LCV

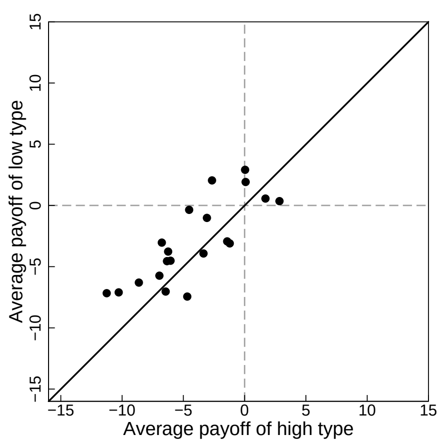

(d) HPV

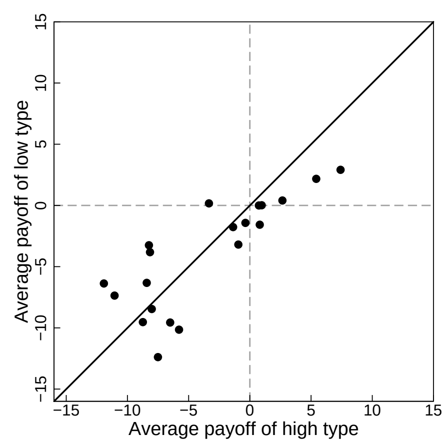

(e) HCV

Figure 9: Scatterplots of observed payoffs conditional on low and high types, by treatment. Each point represents one pair of bidders.

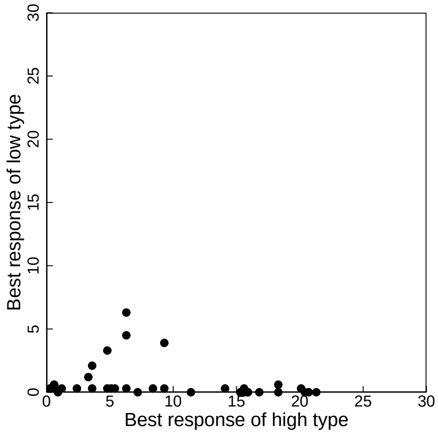

(a) IPV

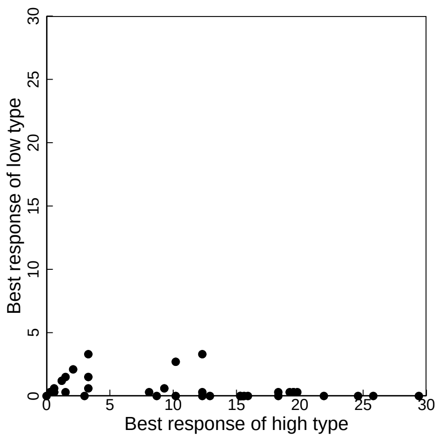

(b) LPV

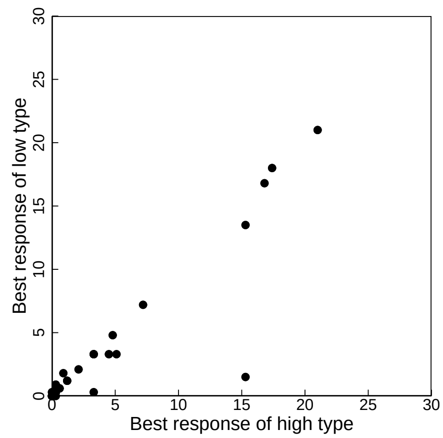

(c) LCV

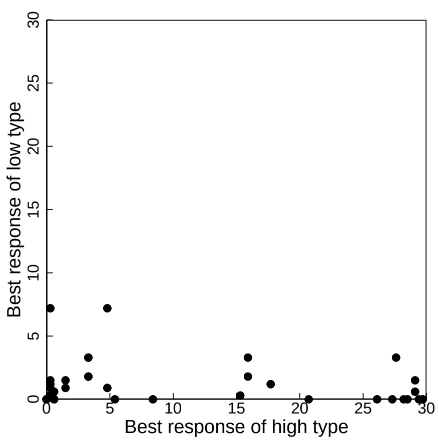

(d) HPV

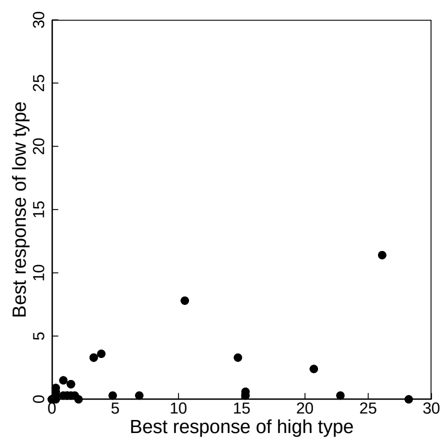

(e) HCV

Figure 10: Scatterplots of best response bids conditional on low and high types, by treatment. Each point represents one bidder.

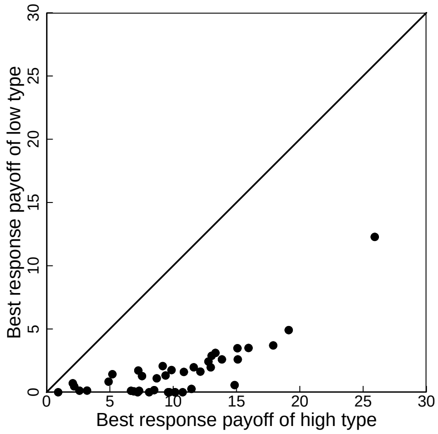

(a) IPV

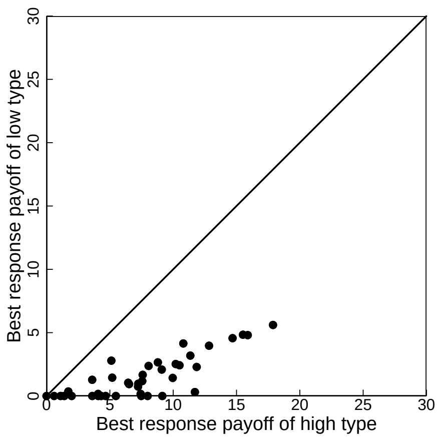

(b) LPV

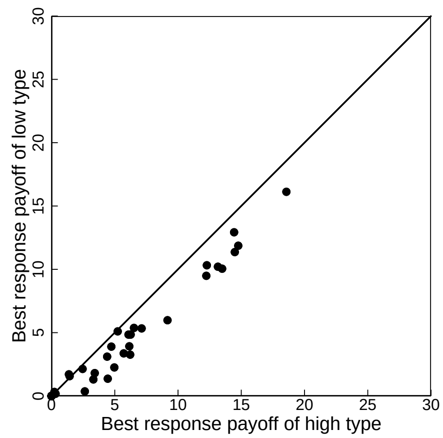

(c) LCV

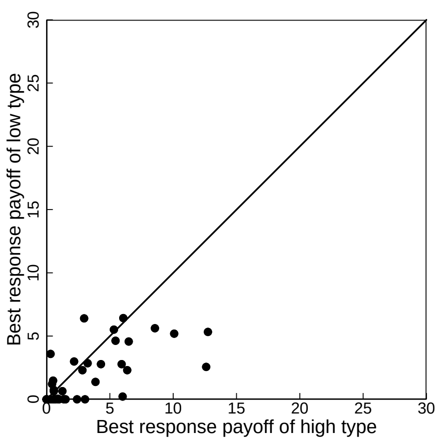

(d) HPV

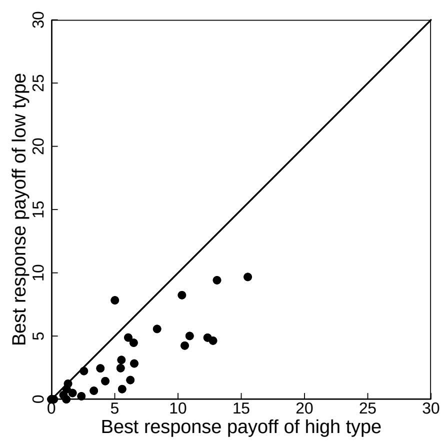

(e) HCV

Figure 11: Scatterplots of best response payoffs conditional on low and high types, by treatment. Each point represents one bidder.

To address the question of whether being type t_H is "good news," we compute for each participant the expected payoff of the best response bids $(b_i^\star(t_H), b_i^\star(t_L))$, which we write $(\pi_i^\star(t_H), \pi_i^\star(t_L))$. Figure 11 presents these payoffs, with each participant corresponding to a point plotted at coordinates $(\pi_i^\star(t_H), \pi_i^\star(t_L))$. If receiving type t_H is indeed good news, we expect to see these best-response payoffs clustered below the 45-degree line shown in the figure.

Result 6. *Measured by the opportunities for best-response payoffs, type t_H is good news relative to type t_L in private values treatments when the monotonicity condition is satisfied. In common values treatments, best-response payoffs are almost equally attractive in LCV, opposite the prediction of monotonicity.*

Support. For IPV and LPV, Figure 11 shows that the best-response payoffs available conditional on t_H are generally much larger than those available to t_L. For HPV, best response payoffs are lower, and cluster more closely to the 45-degree line; in this treatment, receiving type t_H is not significantly better news for payoffs, again as predicted by monotonicity. However, for common values the data do not support the prediction. In LCV, the payoffs available to type t_L are often substantially greater than zero, and almost as large as those available to t_H, whereas in HCV having type t_H is better news relative to t_L; this is the opposite of the prediction of the monotonicity condition.

Table 5 quantifies the graphical observations. For each pair, we take the average of the differences $(\pi_i^\star(t_H), \pi_i^\star(t_L))$ for the two participants in the pair to form an independent pair-level measure, and report summary statistics of this measure in the Table. In IPV, in the average pair receiving type t_H would lead to best-response earnings $Q8.53 \approx \$1.12$ higher than receiving t_L; this premium disappears in HPV. The premium for t_H is small and roughly of the same magnitude for LCV versus HCV, with the mean premium higher on average in HCV in contrast to the equilibrium prediction. □

How well do participants do against this best-response benchmark? The frequency of bids at or near zero could be a naive response to the game, or, given the generally aggressive bidding observed in most pairs, it could be a rational approximate best response to what the other participant in the pair is doing. To investigate we construct a measure of *response quality* as follows. Take some participant i and their co-participant j in the same pair. We take the empirical distribution of bids of j, and, fixing a type t_k for i, compute i's best-response and worst-response payoffs $\overline{\pi}_k$ and $\underline{\pi}_k$ conditional on j's actual bidding behavior. We then compute the expected payoff π_k from the actual bids submitted by i when they had type t_k. The measure of response quality for bidder i when type t_k is then $Q_k = \frac{\pi_k - \underline{\pi}_k}{\overline{\pi}_k - \underline{\pi}_k}$. This normalizes the payoffs onto the $[0, 1]$ interval,

Quantity	Statistic	PV			CV	
		IPV $p = 0.5$	LPV $p = 0.6$	HPV $p = 0.9$	LCV $p = 0.6$	HCV $p = 0.9$
Response quality	Mean over pairs	0.605	0.591	0.578	0.510	0.537
	SD over pairs	(0.070)	(0.080)	(0.083)	(0.149)	(0.127)
	Median over pairs	0.605	0.624	0.586	0.545	0.558
Quality difference	Mean over pairs	0.158	0.163	0.105	0.172	0.110
	SD over pairs	(0.114)	(0.166)	(0.073)	(0.141)	(0.110)
	Median over pairs	0.147	0.905	0.100	0.143	0.072

Table 6: Analysis of response quality measures, by treatment. Response quality refers to average response quality between participants in a pair. Quality difference refers to the difference of response quality within a pair.

with $Q_k = 1$ meaning the participant always chose the best-response bid.[10,11] We then measure i's overall response quality as the average of the response qualities, $Q = (Q_H + Q_L)/2$.

Figure 12 gives scatterplots of the response quality of each participant, plotted against the response quality of their co-participant in their pair; as the plot contains points for both participants in each pair, it is symmetric around the 45-degree line. With PV, responses cluster close to the 45-degree line, indicating that the response quality in each pair is similar for both participants; it is generally not the case that one participant is playing naively while the other best responds. Response qualities are slightly more dispersed in HPV than in IPV or LPV. On the other hand, with CV, response qualities are much more heterogeneous. There are more pairs far from the 45-degree line, indicating one participant in the pair is doing a much better job of best-responding, as well as more dispersed along the diagonal, meaning that some participants are quite far from best-responding in payoff terms.

Result 7. *Response quality is higher in PV than CV, especially in IPV and LPV. Response quality is most heterogeneous in LCV. Overall response qualities are fairly similar within pairs, suggesting few participants are able to exploit systematically their co-participant's strategies.*

Support. Table 6 summarizes two measures of response quality at the pair level. The response quality for a pair is measured as the average of the response quality measures of the members of

[10]The fact that the bid is always paid means there is always substantial variation in the payoffs as a function of bid, and so the range of best-response versus worst-response payoffs is similar for all participants.

[11]Another approach would be to fit a logit noise parameter to the choices, in which case a smaller noise parameter would correspond to more frequent choices of bids giving higher payoffs. The logit model cannot accommodate the case in which the average payoff attained by the participant is less than that which would be attained by uniform randomization; as is seen especially in LCV and HCV, there are participants who do manage to do worse than uniform randomization.

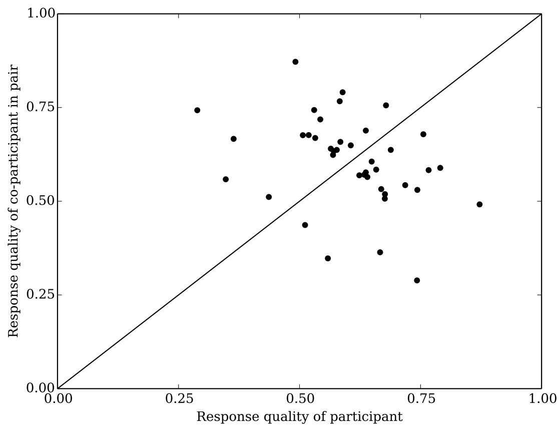

(a) IPV

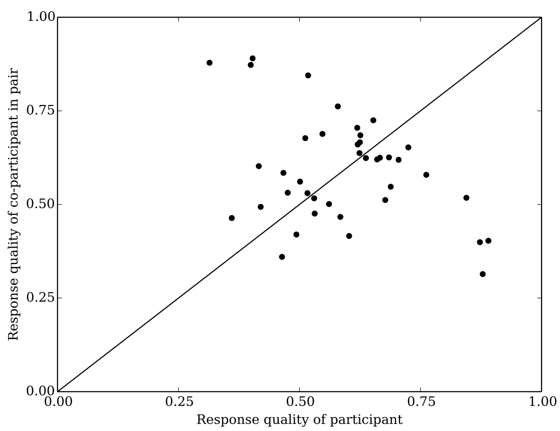

(b) LPV

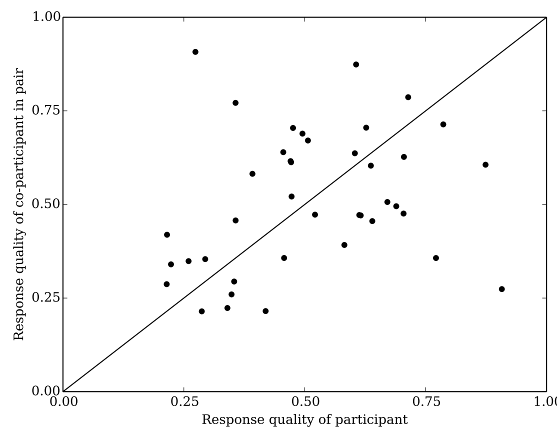

(c) LCV

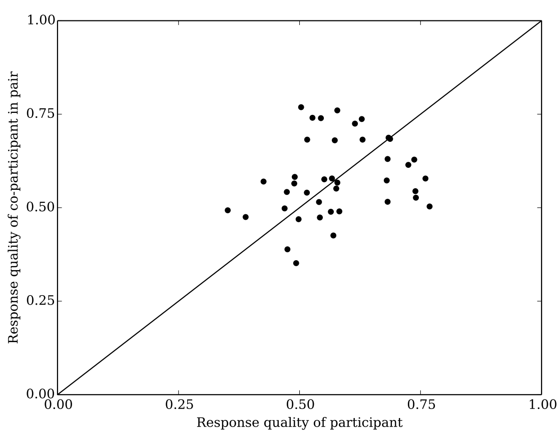

(d) HPV

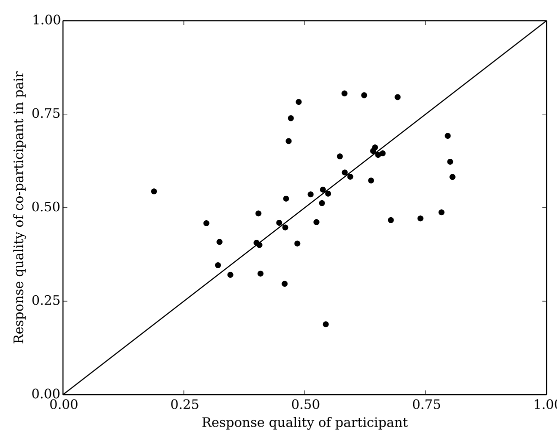

(e) HCV

Figure 12: Scatterplots of response quality, by treatment. Each point represents one participant, and plots their response quality against that of the co-player in the pair; therefore, the plot is symmetric around the 45-degree line.

the pair. Average response quality is highest in IPV at 0.605, and lowest in LCV at 0.510. Response quality measures are also substantially more dispersed in CV than PV.

Heterogeneity within a pair is measured as the difference between the response qualities of the two memvers of the pair. Within-pair heterogeneity is largest in LCV, and is generally larger in CV than PV. Overall neither the scatterplots nor the summary statistics indicate that any participants in any treatment are able to exploit systematically suboptimal play by their counterparts. □

5 Conclusion

Contests are distinguished from, for example, winner-pay auctions by the conditions under which information can be considered unambiguously good news. In winner-pay auctions, affiliation of signals and values is sufficient to ensure that having a "higher" type is definitely good news. The fact that bids in contests are irrevocably sunk even if the contest is lost creates a strategic complexity when types are highly correlated.

We find that the strategic intuition embodied in the monotonicity condition assumed by Krishna and Morgan (1997) and Siegel (2014) is reflected behaviorally in private-values contests: when values are highly correlated, competition is indeed stiff, while higher-value bidders do relatively well when values are uncorrelated or weakly correlated. However, with common values, the implications of the monotonicity condition for the effect of accuracy of information are not observed in our experiment. Theory predicts that bidders with a high signal should do relatively well when information is inaccurate, but the effect of the accuracy of information on behavior is at best mixed.

These results provide a new dimension to the existing literature on behavior in common-value settings. It is generally known that participants in experiments are prone to biases when reasoning about games with common-value elements. Our experiment contributes to this discussion by varying the accuracy of information about the common value, and focusing on the implications of the accuracy on qualitative characteristics of behavior.

A theme of the theoretical analysis of all-pay auctions with private types in Rentschler and Turocy (2016) is the link between the monotonicity condition and complexity. When the matrix ψ satisfies the monotonicity condition, there is a unique equilibrium which has a simple, easy-to-construct form. In the symmetric case, in equilibrium each type "competes" only with the same type of the other bidder. Absent monotonicity, the strategic calculus becomes more complex, as optimal bidding requires consideration of the behavior of multiple possible type realizations of the other player. The IPV and LPV parameterizations are simple in this sense, and indeed those are the ones for which the qualitative predictions of the theory are observed in the data. In HPV, types t_H optimally should also be in direct competition sometimes with an opponent of type t_L; this is a more complex situation, and we do observe lower and more dispersed response qualities in this setting.

The monotonicity condition on ψ is elegant in that it does not distinguish between private, common, or interdependent valuation structures. It does however rely crucially on conditional probabilities and conditional expectations, two tasks at which humans are not generally strong. Therefore, hidden in the monotonicity condition is another way in which parameterizations can differ in complexity; the conditioning in CV is substantially more challenging than in PV. Our results suggest that it is this complexity which dominates behavior: in LCV, participants of type t_H don't realize that it is quite likely that they are, in fact, the prize's only admirer.

References

Susan Athey. Single crossing properties and the existence of pure strategy equilibria in games of incomplete information. *Econometrica*, 69(4):861–889, 2001.

Christiane Ernst and Christian Thöni. Bimodal bidding in experimental all-pay auctions. *Games*, 4(4): 608–623, 2013.

Brit Grosskopf, Lucas Rentschler, and Rajiv Sarin. Asymmetric information in contests: theory and experiments. Technical report, Working Paper, Texas A&M University, 2014.

Sergiu Hart. Discrete Colonel Blotto and General Lotto games. *International Journal of Game Theory*, 36: 441–460, 2008. doi: 10.1007/s00182-007-0099-9.

Hannah Hörisch and Oliver Kirchkamp. Less fighting than expected. *Public Choice*, 144(1):347–367, 2010.

Kyle Hyndman, Erkut Y. Ozbay, and Pacharasut Sujarittanonta. Rent seeking with regretful agents: Theory and experiment. *Journal of Economic Behavior & Organization*, 84(3):866–878, 2012.

Vijay Krishna and John Morgan. An analysis of the war of attrition and the all-pay auction. *Journal of Economic Theory*, 72:343–362, 1997.

Paul R. Milgrom and Robert J. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50: 1089–1122, 1982.

Wieland Müller and Andrew Schotter. Workaholics and dropouts in organizations. *Journal of the European Economic Association*, 8(4):717–743, 2010.

Charles Noussair and Jonathon Silver. Behavior in all-pay auctions with incomplete information. *Games and Economic Behavior*, 55(1):189–206, 2006.

Lucas Rentschler and Theodore L. Turocy. All-pay auctions with interdependent valuations, including the highly competitive case. *Journal of Economic Theory*, 2016.

Roman Sheremeta. Overbidding and heterogeneous behavior in contest experiments. *Journal of Economic Surveys*, 27(3):491–514, 2013.

Ron Siegel. Asymmetric all-pay auctions with interdependent valuations. *Journal of Economic Theory*, 153 (C):684–702, 2014.

Theodore L. Turocy and Timothy N. Cason. Bidding in first-price and second-price affiliated-values auctions: A laboratory experiment. Working paper, 2015.

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.

Ruqu Wang. Common-value auctions with discrete private information. *Journal of Economic Theory*, 54 (2):429–447, 1991.