Reservation Values and Regret in Laboratory First-Price Auctions: Context and Bidding Behavior*

Theodore L. Turocy† Elizabeth Watson‡

August 1, 2011
Forthcoming, *Southern Economic Journal*

Abstract

Recent papers hypothesize that an asymmetry in regret motivates aggressive bidding in laboratory first-price auctions. Subjects emphasize potential earnings foregone from being outbid. Proposed motivators of this asymmetry include the one-to-one relationship in the auction between winning and positive earnings and the ex-post knowledge that bidders who do not win the auction know they earned less than the winning bidder. We design a novel implementation of the first-price auction environment in which these characteristics are not present, while leaving unchanged the expected-earnings maximizing bidding strategy against any fixed beliefs about the bidding behavior of others. Bidding is significantly less aggressive in this treatment. These findings support the hypothesis that aggressive bidding is motivated in part by features of the protocol for incentivizing subjects which are not essential to the auction environment.

Keywords: first-price auctions, regret, framing, methodology of experiments
JEL classifications: D44, C90, C91

1 Introduction

> "If you've never been surprised by the effect that a small change in instructions can have, then you've never tried it." – Vernon Smith, keynote plenary session, 2007 World Meeting of the Economic Science Association, Rome, Italy, June 2007.

*Funding for the reported experimental sessions was provided by the Department of Economics at Texas A&M University. We thank the Economic Research Laboratory at Texas A&M for use of facilities and logistical support in running these sessions. We also thank Glenn Harrison, Charlie Holt, Robert Sugden, and participants at several Economic Science Association meetings for helpful comments.

†School of Economics and Centre for Behavioural and Experimental Social Science, University of East Anglia, Norwich, Norfolk NR4 7TJ, United Kingdom; email: T.Turocy@uea.ac.uk

‡Department of Economics, Trinity University, San Antonio TX 78212, United States; email: ewatson1@trinity.edu

A long-standing puzzle in laboratory first-price private-values auctions is that subjects bid significantly more aggressively than predicted by the risk-neutral Nash equilibrium. The survey of Kagel (1995) outlines the history of this result. The financial consequences to subjects of this aggressive bidding are substantial relative to the scale of incentives. For instance, in Turocy, Watson, and Battalio (2007), subjects in first-price private-values auctions averaged earnings around \$13.00 for participating in 60 auction periods; if a typical subject had unilaterally switched to the risk-neutral equilibrium strategy, she would have increased her earnings by \$8.00.

This stylized fact warrants the close attention it has received in the literature due to the broader roles of laboratory experiments in auctions. Auction experiments have been used to refine auction models and theories of bidding (Cox et al 1982), and have informed the design of auction mechanisms in the field (for example, Roth 2002). In order to use the results of laboratory experiments to draw inferences about the determinants of bidding behavior in real-world auctions, we must be able to assume that subjects focus on the same strategic considerations that we believe agents in the field entertain. In the standard theory of bidding in first-price private-values auctions (Vickrey 1961, Milgrom and Weber 1982), the strategic consideration faced by a bidder is a price-probability tradeoff. A higher bid increases the probability of winning the auction, but decreases the consumer surplus the bidder gains when she wins, because she pays a higher price.

Some recent experiments, including Engelbrecht-Wiggans and Katok (2007, 2008, 2009), Filiz-Ozbay and Ozbay (2007), and Ockenfels and Selten (2005), investigate the hypothesis that aggressive bidding in laboratory first-price auctions is motivated by a difference in how subjects anticipate or react to the outcomes of winning versus not winning an auction. This line of inquiry is based on the observation that the bid which is optimal ex ante will not in general be the optimal bid ex post, after the results of the auction are known. If the bidder wins an auction by submitting a bid strictly higher than any other bid, she may realize that she would have been able to win the auction had she submitted a lower bid, and therefore could have increased her earnings; this is "winner regret." Conversely, if the bidder loses the auction, but the winning bid turns out to be less than her value, she could have bid higher and made positive earnings instead of zero earnings; this is "loser regret." Theoretical results in the literature on regret in auctions, dating to Engelbrecht-Wiggans' (1989) analysis of optimal bidding with regret, predict that higher bids result when loser regret is relatively more salient to subjects than winner regret.

In this paper we examine how the exposition of the laboratory auction environment affects the strategic considerations faced by a bidder and might lead to the greater salience of loser regret implied by the results in the literature on regret in auctions. In most first-price auction experiments to date, subjects are instructed that they have a privately-known idiosyncratic "resale value" for the object being auctioned. The winning bidder purchases the object at the price she bid. She then sells it back to the experimenter for her resale value and earns the difference between her resale value and her bid. The other bidders do not purchase the object, have nothing to re-sell to the experimenter, and thus have earnings of zero. As this amounts to the subject earning the amount of her profit from her activity in the auction, we refer to this as the "profit frame." This design of incentives has three interrelated properties which may plausibly influence how a subject processes her decision problem.

1. The only subject who makes positive earnings in an auction period is the winning bidder.[1]

[1] Engelbrecht-Wiggans and Katok (2008) report that executive MBA students who participate in informal classroom auctions immediately notice this feature.

2. The winning bidder is the only participant who appears to engage successfully in an economic activity in the period.

3. The winning bidder enjoys what Ockenfels and Selten call "advantageous relative standing," in that by winning the auction, she earns more in that period than the participants against whom she was bidding.

We introduce an alternative protocol in which all of these considerations are absent and show bidding behavior is significantly less aggressive. The "surplus frame" treatment places the auction in a broader economic context. Subjects have privately-known idiosyncratic "outside prices" for a close substitute of the object being auctioned. The winning bidder purchases the object in the auction for the price she bid; the other bidders purchase outside the auction at their outside price. All bidders earn money equal to their consumer surplus from the purchase.[2] All bidders make a purchase and make positive earnings in each period. Therefore, the ex-post relative standing of bidders can no longer be inferred solely from the result of the auction. In fact, in the surplus frame, the winning bidder is predicted to be the bidder with the worst relative standing, conditional on the realized outside prices, prior to bidding.

The paper proceeds as follows. Section 2 reviews the standard Bayesian game model of the first-price auction and outlines the predictions of the regret-based models. Section 3 describes the experimental protocol, with particular detail on the surplus frame manipulation. Section 4 reports the results. Section 5 evaluates the results within the greater context of the literature to date and discusses the implications of the results for using the laboratory as a testbed for auction-like environments.

2 Theory

2.1 The standard first-price auction model

The standard game-theoretic model of a first-price auction for a single, indivisible object is a Bayesian game. There are $N \geq 2$ bidders. Each bidder i has a type x_i, which represents the largest amount she is willing to pay for the object in the auction. The bidder knows her type prior to bidding. We focus on the case in which types are drawn independently for each bidder from the uniform distribution over some support $[\underline{x}, \overline{x}]$. Bids are submitted simultaneously, with each bidder i submitting a bid $b_i \geq 0$, without knowledge of any other bidder's realized type. The object is purchased by the bidder submitting the highest bid and the winning bidder pays the amount of her bid. In the event of a tie, one of the bidders tied at the highest bid is selected at random as the winner.

Under the assumptions that bidders are risk neutral[3] and care only about expected earnings,

[2]Therefore, the resale value and outside price are isomorphic, insofar as a bidder with resale value x in the profit frame has the same contingent incentives as a bidder with the same outside price x in the surplus frame. Both the resale value and the outside price play the role of the maximum price a subject is willing to pay for the object in the auction. Thus, they both represent the reservation value of the theoretical model. This interchangeability between resale value and outside price is also noted by Kirchkamp et al (2009).

[3]In Appendix A we extend the analysis to consider risk aversion and show the results of our experiments are inconsistent with risk aversion as a maintained hypothesis. This provides independent support to the rejection of risk aversion in favor of regret by Engelbrecht-Wiggans and Katok (2009).

a bidder with type x_i maximizes her expected utility in the auction by solving the optimization problem

$$\max_{b_i} P(b_i)(x_i - b_i), \tag{1}$$

where $P(b)$ is the probability a bid of b wins the auction, taking into account the expected behavior of other bidders. The symmetric Bayes-Nash equilibrium (Vickrey 1961) involves all bidders adopting the bid function

$$b(x) = \frac{N-1}{N}x. \tag{2}$$

2.2 Regret-based approaches

A robust finding in the literature on the first-price auction in the laboratory is that most subjects bid more aggressively than predicted by the risk-neutral Nash equilibrium. Several theoretical and experimental papers propose to explain this result by assuming that the winner regret and loser regret outcomes have different effects on the process of formulating bids. There are two strands to this literature. First, following Engelbrecht-Wiggans (1989), the model of an expected utility-maximizing bidder is augmented to include (additive) terms incorporating the effects of anticipated regret. These models then maintain the condition that bidders bid in accordance with a strategic (Bayes-Nash) equilibrium assuming the augmented utility functions. A related approach, offered by Ockenfels and Selten (2005), instead takes ideas from learning theory to formulate an equilibrium concept in which the two types of regret balance on average. We review each of these ideas briefly in turn.

Several closely-related specifications of the anticipated regret strategic equilibrium model appear in Engelbrecht-Wiggans (1989), Engelbrecht-Wiggans and Katok (2007, 2008, 2009), and Filiz-Ozbay and Ozbay (2007). These models augment the formulation in (1) by adding penalties to the bidder's objective function to model anticipated regret. Taking the perspective of bidder i, consider an outcome of the auction in which the maximum among bids submitted by bidders other than bidder i is w, and m other bidders submitted a bid of w. In this contingency, the utility to a bidder with type x_i who bids b_i is given by

$$u(x_i, b_i, m, w) = \begin{cases} x_i - b_i - R_W(b_i - w) & \text{if } b_i > w \\ \frac{1}{m+1}(x_i - b_i) - \frac{m}{m+1} R_L(x_i - w) & \text{if } b_i = w \\ -R_L(x_i - w) & \text{if } b_i < w. \end{cases} \tag{3}$$

The function $R_W(\cdot)$ captures *winner regret*; ex post, the winning bidder finds out that with a lower bid she could still have won the auction, but paid a lower price. The function $R_L(\cdot)$ captures *loser regret*; ex post, a losing bidder finds out that with a higher bid she could have won the auction profitably at a price below her reservation value. Both functions are assumed to be non-negative and non-decreasing in their argument, and satisfy $R_W(z) = 0$ and $R_L(z) = 0$ for all $z \leq 0$. The original specifications of Engelbrecht-Wiggans, and later of Engelbrecht-Wiggans and Katok, assume linear forms for the regret functions, $R_W(z) = \alpha_W z$ and $R_L(z) = \alpha_L z$, while Filiz-Ozbay and Ozbay consider a more general form. In both versions, equilibrium bidding has intuitive comparative statics: increasing the weight placed on loser regret makes equilibrium bidding more aggressive and increasing the weight placed on winner regret makes equilibrium bidding less aggressive.

Ockenfels and Selten (2005) also adopt the notion of differential effects of regret in proposing the application of impulse balance equilibrium. Instead of being formulated within a strategic equilibrium framework, impulse balance equilibrium is motivated by ideas from learning models, in which bidders adjust their bidding behavior based on the outcomes of previous auctions. The case of loser regret is treated as an instance of *upward impulse*; because the bidder missed out on an opportunity to win the auction profitably, she is impelled to bid more aggressively in future. The case of winner regret is defined as an instance of *downward impulse*; a bidder who has won an auction at a price strictly higher than her opponents' bids is impelled to bid less aggressively in the future. In a weighted impulse balance equilibrium, the expected upward impulse is equal to λ times the expected downward impulse, where λ is a parameter representing the relative salience of upward and downward impulses. The term "equilibrium" in impulse balance equilibrium, then, is not the standard notion of strategic equilibrium, but rather is akin to the equilibrium of a physical system, in which the countervailing forces of the two types of regret exactly balance out. Ockenfels and Selten show that as λ decreases, which implies greater weight placed on the loser regret outcome, bidding in impulse balance equilibrium becomes more aggressive.

The key observation is that both classes of model adopt the notion that winner regret and loser regret enter the bidder's bid formulation process differently. Both models also make the same qualitative comparative statics prediction, that as the weight on loser regret increases, bids become more aggressive. We view these two notions of equilibrium as complementary perspectives, and leave for future research the question of designing a test which might be able to distinguish the predictions of strategic versus impulse balance equilibrium.

Previous experiments have generally used the structure of the ex-post feedback on results of an auction to probe the consequences of the distinction between winner and loser regret. In contrast, our experimental design considers how asymmetric regret may arise based on the ex-ante exposition of the auction environment.

2.3 The surplus frame interpretation

In the experiments reported in this paper, we introduce the surplus frame auction environment. Again, there is a single indivisible object for auction. There are $N \geq 2$ bidders. Each bidder has a maximum value for the object $v \geq 1$; this value is the same across all bidders. Each bidder i is aware of an opportunity to purchase a close substitute to the object at a price x_i, which she knows in advance of the auction game. These "outside prices" define the bidders' types and are drawn independently from the uniform distribution over some interval $[\underline{x}, \overline{x}]$. Bids are submitted simultaneously, with each bidder i submitting a non-negative bid b_i without knowledge of any other bidder's realized type. The object being auctioned is purchased by the bidder submitting the highest bid, and the winning bidder pays the amount of her bid. In the event of a tie, one of the bidders tied at the highest bid is selected at random as the winner. Losing bidders purchase the close substitute at their respective outside prices x_i.

Under the assumptions that bidders are risk neutral and care only about expected earnings, a bidder maximizes her expected utility by solving the optimization problem

$$\max_{b_i} P\left(b_i\right)\left(v - b_i\right) + \left(1 - P\left(b_i\right)\right)\left(v - x_i\right), \tag{4}$$

where $P(b)$ is the probability a bid of b wins the auction, taking into account the expected behavior of other bidders. This rearranges to

$$(v - x_i) + \max_{b_i} P(b_i)(x_i - b_i). \tag{5}$$

The expression (5) from the surplus frame is identical to (1), except for the addition of a type-dependent constant. Therefore, the Bayes-Nash equilibrium (2) is unchanged. A fortiori, for *any* fixed beliefs about the behavior of other bidders, encoded in $P(b)$, the expected-earnings maximizing bid is the same in both frames.

The predictions of the regret and impulse balance models are also unchanged in the surplus frame, holding fixed the parameters of those models. The regret functions $R_W(\cdot)$ and $R_L(\cdot)$ take expected absolute amounts of foregone earnings as arguments; if these are independent of frame, then it follows immediately that the anticipated regret equilibrium is unchanged. Similarly, the addition of an additive constant to payoffs leaves impulse balance equilibrium unchanged, if the parameter λ is assumed independent of the frame.

However, the surplus frame design is intended to manipulate the proposed qualitative motivators of asymmetric regret. In the surplus frame, subjects earn money even when they do not win the auction. More than simply moving the losing earnings amount away from zero, the type-specific additional payoff in (5) is inversely related to the outside price. This feature addresses the hypothesis of Ockenfels and Selten (2005) that "differences in relative standing" with regards to earnings motivate greater weight on loser regret. In the profit frame, only the bidder who wins has positive earnings. While she might have been able to earn more in hindsight, she nevertheless has an advantageous relative standing towards all the other bidders, insofar as she received positive earnings while all other bidders earned zero. In contrast, a bidder suffering loser regret realizes, had she bid differently, she could have been the one to take away positive earnings instead of zero earnings. In the surplus frame, winning the auction does not equate with advantageous relative standing. A bidder with an outside price of $\$5.10$ who wins the auction with a bid of $\$4.00$ earns less than a bidder who had an outside price of $\$3.00$ but who did not win the auction. Roughly speaking, in the profit frame, having a high realization of x_i is *good* news for the subject; in the surplus frame, it is *bad* news.

Therefore, if the close relationship between winning and positive earnings, and differences in relative standing, jointly influence the weight bidders place on loser regret, both the anticipated regret and weighted impulse balance models predict less aggressive equilibrium bidding under the surplus frame than the profit frame.

3 Design

The design extends the protocol of Turocy, Watson, and Battalio (2007) (TWB). We report results on a total of 6 experimental sessions, three using the profit frame and three the surplus frame. Each session consisted of 9 subjects and lasted 60 periods. In each period, the bidders were divided into three markets, with three bidders each. Subject assignments to markets were independent across periods, and types were uniformly distributed and independent across periods and subjects. The same assignments to markets and realized types were used for all sessions. Each bidder received a type each period from the set {\$0.15, \$0.30, ..., \$5.85, \$6.00}. After learning their types for the period, subjects simultaneously chose a bid from the set $\{\$0.10, \$0.20, ..., \$6.10, \$6.20\}$.[4] The bidder submitting the highest bid purchased the object in the auction at a price equal to her bid. Ties were resolved by choosing one of the tied bidders at random.

The design manipulates the presentation of the type. In the control sessions using the profit frame, the type is presented as a "resale value" at which the subject can "re-sell" the object to the experimenter if she wins the auction. The part of the instructions informing subjects of their earnings calculation read

> Your Earnings for a period will depend on whether you purchase the commodity in your market, and on the Market Price.
>
> If you purchase a unit of the commodity, your earnings for that period will be calculated according to the equation
>
> Your Earnings = Resale Value - Market Price
>
> If you do not purchase a unit of the commodity, then your earnings for that period will be zero.

In sessions using the surplus frame, the type is motivated as an "outside price" at which the subject can purchase an identical object outside of the auction if she does not win the auction. The instructions regarding earnings calculations instead used this language:

> You will purchase exactly one unit of the commodity each period. If you purchase the unit of the commodity in the market, your earnings for that period will be calculated as
>
> Your Earnings = \$6.20 - Market Price
>
> If you do not purchase the unit of the commodity in the market, then you will purchase a unit outside the market at your Outside Price. Your Earnings for the period are then computed as
>
> Your Earnings = \$6.20 - Outside Price

The remainder of the instructions was identical, except for substituting terminology where required.

Also identical, up to changes in terminology, was the graphical computer interface the subjects used to receive information and make their decisions. In addition to the display of the current auction period, the screen contained a record sheet reporting the results of the last 25 periods, with scroll buttons available to view earlier periods once filled. In the profit frame sessions, subjects were paid their total earnings from all 60 periods, plus a \$5.00 initial balance; the record sheet kept

[4] In general, in passing to a discretized version of a continuous model, the properties of the continuous equilibrium need not be preserved. The discretization of reservation values in increments of \$0.15 and bids in increments of \$0.10 was selected so that it is an equilibrium, if bidders are all risk-neutral, to bid two-thirds of the private value in both frames.

a running total of earnings, with the \$5.00 balance already included at the start of the session. In the surplus frame sessions, subject earnings are positive in all 60 periods. As such, paying subjects the sum of their earnings from all 60 periods would result in much higher expected earnings than in the profit frame where most subjects have 0 earnings in most periods. In order to maintain the same level of expected earnings across sessions, assuming the same bidding behavior, subjects in the surplus frame sessions were paid their earnings from 7 of the 60 periods, with no initial balance. This was announced in the instructions for the session. The periods which were paid were selected after all 60 periods were completed by physically drawing numbered chips from a cup in front of all subjects.

Each session consisted of 9 subjects recruited from the undergraduate student body at Texas A&M University. No subject participated in more than one session, and no subject had previously participated in any auction experiment. All interaction among the subjects was mediated via computer. All matching and bidding was done anonymously; no ID numbers or other identifying information was made known to the subjects. At the end of each period, subjects only found out the highest bid in their market; no information about other bids was revealed.[5]

4 Results

4.1 Auction revenue

We begin the analysis of the results at the market level, and examine the percentage of the possible gains from exchange which accrue to the seller in both frames. We follow Ockenfels and Selten (2005) in organizing the analysis around relative bids, i.e., bids expressed as a percentage of the bidder's type. In doing so, when we state results about the treatment effect of the surplus frame on revenue, we automatiaclly control for the varying realizations of the maximum type in each auction. Further, the equilibrium bid function with both risk-neutral and CRRA bidders is linear, as is the specification of impulse balance equilibrium made by Ockenfels and Selten. Finally, as Ockenfels and Selten note, approximately linear bid functions organize the behavior of many subjects in previous experiments. These combine to make relative bids the natural quantity for analysis.

Table 1 summarizes auction performance at the session level. The risk-neutral equilibrium prediction for the mean winning bid is \$3.025.[6] Mean winning bids exceed this prediction in all sessions. The overall mean winning bid in the profit frame is \$3.773, and in the surplus frame \$3.399. As a point estimate, the surplus frame explains almost exactly 50% of the excess revenue in the profit frame relative to the risk-neutral Nash baseline. An equivalent expression of these data is found by examining mean relative winning bids. The mean relative winning bid is defined as the maximum bid in an auction, divided by the maximum realized type, and, therefore, captures the percentage of the possible gains from exchange which accrue to the seller. Here, the averages in the profit frame (0.851) and the surplus frame (0.769) both exceed the risk-neutral baseline of $\frac{2}{3}$.

[5]The results of Filiz-Ozbay and Ozbay (2007) and Ockenfels and Selten (2005) both indicate that this feedback pattern tends to increase bids. We maintain this pattern to give the surplus frame a more strenuous test relative to the risk-neutral baseline.

[6]In the uniform continuous model, the mean winning bid for values distributed uniformly on $[\$0.00, \$6.00]$ is \$3.00. The prediction of \$3.025 arises in our discretized setting because we do not assign the value of 0 to any bidders.

Figure 1 takes a time-series approach to the data by plotting ten-period moving averages of the relative winning bid for each of the six sessions. This leverages the direct comparability of sessions due to the use of the same realizations of types and of matching of subjects into groups. The plot provides striking visual evidence of a treatment effect: the time series for the sessions using the surplus frame lie everywhere below the series for sessions using the profit frame.

Finally, at the individual auction level, Figure 2 displays a scatterplot of all auction outcomes as measured by the mean relative winning bid. Three main features are evident:

1. In both treatments, a majority of markets exceed the risk-neutral baseline.

2. Auctions conducted using the surplus frame generally result in lower relative winning bids.

3. In both treatments, there is evidence that the relative winning bid decreases when the maximum type in the auction is larger, with the effect being more pronounced in the surplus frame.

We formalize these graphical observations using a regression model to predict relative winning bids as a function of market characteristics. Let s index the session, t index the period number (1 to 60), and i index the market. Let S_{its} be a dummy variable equaling one when the session used the surplus frame, and P_{its} be a dummy variable equaling one for the profit frame. Let b_{its}^{MAX} be the winning bid in auction its, and x_{its}^{MAX} be the maximum type in the auction. We estimate the model

$$\frac{b_{its}^{MAX}}{x_{its}^{MAX}} = \alpha_0 P_{its} + \alpha_1 S_{its} + \beta_1 t P_{its} + \beta_2 t S_{its} + \beta_3 x_{its}^{MAX} P_{its} + \beta_4 x_{its}^{MAX} S_{its} + \varepsilon_{its}. \tag{6}$$

Table 2 reports the coefficient estimates and standard errors on the coefficients. We cluster errors at the session level. The signs of the point estimates of all coefficients confirm the graphical impressions. We test the null hypothesis that there is no treatment effect using the joint hypothesis $H_0 : \alpha_0 = \alpha_1, \beta_3 = \beta_4$. The null hypothesis is rejected in a two-sided test (p-value 0.0012), providing evidence for a treatment effect. The lower revenues in the surplus frame are significant both statistically and economically.

Focusing specifically on the time trend, while the point estimates of the time trend are negative in both frames, only the coefficient on the time trend in the profit frame is significant. To confirm this result is robust to the choice of specification, we also carry out a nonparametric test (Cuzick 1985) for trend and incorporate all bids rather than solely the winning bids. For each session s, we compute for each period t the median relative bid among all nine bidders in that period in that session. This generates an independent time series for each session, which form the data for the test. The null hypothesis is that there is no trend over time. In the profit frame, we can reject the null hypothesis (p-value 0.010), but in the surplus frame we cannot (p-value 0.514). This is in good agreement with the results obtained via the linear regression.

Our time trend result complements a finding in Engelbrecht-Wiggans and Katok (2008). Their design, which manipulates the feedback the subjects receive, employs "winner regret" and "loser regret" treatments. They find a significant decrease in bids over time in their winner regret treatments, but not under loser regret. They propose as an explanation that under the profit frame, subjects anticipate loser regret, in part perhaps because they recognize straightaway that the only way to make positive earnings is to win, but only incorporate winner regret through experience.

This explanation is consistent with our results comparing profit versus surplus frame, insofar as the surplus frame breaks the direct link between winning and positive earnings, and casts the bidding problem as one of searching for a better price.

4.2 Individual bidding behavior

Since revenues are systematically lower in sessions using the surplus frame, bidding must be less aggressive in aggregate. We turn to analyzing the underlying bidding structure that generates the revenue results.

Because revenues are determined by the highest bidder in an auction, revenue analysis may reflect disproportionately the most aggressive bidders in the population. We therefore complement the revenue analysis by summarizing typical bidding behavior. Selten and Buchta (1999) provide evidence that subjects take a holistic perspective in determining their bidding strategy. Therefore, for each bidder i in each session s, we take the type and bid pairs for each of the 60 periods and nonparametrically estimate a bid function $\hat{b}_{is}$ using lowess regression. Using these estimated bid functions, we then compute for each session the median of the nine bid functions from that session for each of the possible bidder types.[7] Figure 3 displays these median bid functions for each of the six sessions. Consistent with the results on revenue, bidding is visibly less aggressive for high realizations of the private value in the surplus frame. For low realizations of the private value, bids are comparable between treatments. On the whole, the median bid in the profit frame is slightly concave, while in the surplus frame the concavity is much more pronounced.

We further break out the analysis of bid function shapes by individual bidder. As a rough measure[8] of the shape of the bid function, we compute for each bidder

$$\kappa_{is} = \frac{\hat{b}_{is}(\$6.00) - \hat{b}_{is}(\$4.05)}{\$1.95} - \frac{\hat{b}_{is}(\$2.10) - \hat{b}_{is}(\$0.15)}{\$1.95}.$$

That is, κ_{is} is the difference in the estimated slope of the bid function, taken between the upper third of possible types and lower third of possible types. The sign of κ_{is} captures the gross shape of the bid function, with $\kappa_{is} < 0$ corresponding to bidding behavior which is qualitatively concave. Figure 4 displays cumulative distribution functions of κ_{is} for each treatment. In both treatments, a majority of bidders are estimated to have bid functions that are concave to at least some degree. For $\kappa_{is} < 0$, the distribution function curve for the surplus frame lies everywhere to the left of that for the profit frame. The greater concavity in median bids observed in the surplus frame is, therefore, supported by a shift across the whole population towards bidding behavior which is flatter as a function of type when the bidder has a high type. Several bidders in the surplus frame have estimated values of κ_{is} nearing -1, which corresponds, roughly, to bidding approximately equal to the type for low realizations of the type, while bidding approximately a constant amount when the type is high.

[7] By taking the medians of the bid functions, as opposed to the medians of all bids, we weight each subject in the session equally. Each subject faced a different realization of types, and therefore, at any type, we have more observations for some subjects than others. Carrying out this exercise using the median actual bid for each type yields qualitatively and quantitatively similar results, so our conclusions are not sensitive to this choice of methodology.

[8] We experimented with several variations on the measure we present here; the results are not sensitive to our operationalization of concavity.

5 Conclusion

The surplus frame implementation of the first-price auction modifies or removes several characteristics which have been proposed to motivate asymmetry in regret, and therefore increase subjects' aggressiveness in bidding. Consistent with these explanations, bidding is less aggressive, and therefore the implied seller revenues are lower, when using the surplus frame relative to the standard profit frame. In both frames, bidders typically adopt bid functions which are concave in their type, with the concavity being more significant in the surplus frame.[9] These results indicate a potentially greater role for framing in motivating auction environments than has been appreciated to date in the literature.

Independently of our study, Armantier, Holt, and Plott (2010) also explicitly address profit versus surplus framing in auctions. In their setting, subjects, in the role of banks, are bidding to sell distressed (or "toxic") assets to a government agency. As in the recent real-world crisis which it modeled, a key feature of the design was that the subjects did not themselves know the value of the assets they held. This exposed the subjects to the possibility of selling an asset for less than it was actually worth. In describing the reasons for considering both profit and surplus frames, they wrote:

> In particular, paying nothing on shares not sold may provide subjects with an incentive to bid aggressively in order to improve their chances of selling their shares to the government and thereby earn money. On the other hand, because of loss aversion, it may lead subjects to bid conservatively in order to avoid losses. In other words, paying nothing for the shares not sold at the auction could affect behavior, although the direction of the effect is difficult to predict. (p. 27)

In the first-price auction setting where subjects do know their types precisely, subjects can be sure of bidding less than their types, and therefore never face the risk of loss. Therefore, all that remains is the incentive to improve the chances of making positive earnings. Our results provide circumstantial evidence that both the factors Armantier et al were concerned with were at play, and suggest that the presence of loss aversion, which is potentially relevant in their setting but not in ours, explains why we find a significant effect due to frame while Armantier, Holt, and Plott do not.

Kirchkamp et al. (2009) also use an outside option device in a first-price auction experiment. Their profit-frame design incorporates a nonzero outside option payoff in the event the bidder does not win the auction. They show, in a parallel development to our equations (4) and (5), that the bidding problem with a private value and outside option profit reduces to the standard one-dimensional auction model. They find that bidding is even more aggressive, relative to the risk-neutral Nash baseline, when private values and nonzero outside option payoffs are both present. Their results complement ours by demonstrating that our finding of less aggressive bidding is not driven solely by offering a nonzero payoff to losing bidders in the surplus frame.

The surplus frame design also addresses a concern articulated by Harrison, List, and Towe (2007) in the context of a field auction for collectible coins:

[9] The concavity of bids in the surplus frame is inconsistent with CRRA, and the aggressive bidding for low types is inconsistent with any explanation based on risk attitudes; see Appendix A for details.

> We hypothesize that there is a danger that the imposition of an exogenous laboratory control might make it harder, in some settings, to make reliable inferences about field behavior. The reason is that the experimenter might not understand something about the factor being controlled, and might impose it in a way that is inconsistent with the way it arises naturally in the field... (p. 433)

The standard Bayesian game model of a first-price auction posits that bidders have a well-defined maximum willingness to pay. In the profit frame, this is implemented using the resale value mechanism, which establishes a difference in kind between the two outcomes of winning or not winning the auction. In a profit frame experiment, the only way to earn a positive amount of money, or to do better than others in the cohort, is to win. This difference in kind is further emphasized by the necessity of textually differing explanations for how earnings are calculated in the two outcomes. In the surplus frame, subjects always purchase an object, and earnings are always calculated as value minus price paid (i.e. total consumer surplus), irrespective of whether the auction is won or lost. Loosely speaking, the profit frame highlights getting or not getting a special object; the surplus frame highlights participating in an auction to attempt to get a better price.

Neither the surplus frame nor the profit frame is unambiguously better or worse for implementing first-price private values auctions. The appropriate design choice depends on the ultimate aims of the experimental project. The profit frame sets up the object in the auction as being unique; indeed, in a sense, it is the only object available in the subject's simulated universe, since no substitutes are mentioned. The object is presented to the subject surrounded by a halo of stark incentives: to earn any money in the experiment, victory in the auction is essential. This is the appropriate frame when an auction design is being contemplated for selling distinctive items with no close substitutes, for instance a rare van Gogh or an asset crucial to a business plan. However, online auction markets like eBay feature thousands of auctions for items which are relatively common. Auction applications for these more common items may demand a laboratory translation which is less black-and-white, and sets the auction within a broader economic context.

Our results demonstrate that in the standard single-object, private-values case, the choice of profit versus surplus frames does matter, and therefore the experimenter or practitioner has a nontrivial choice to make in designing laboratory experiments intended to shed light on broader-world behavior.

A Best response and equilibrium with risk aversion

In the body of the paper we focus on the risk-neutral Nash equilibrium as a baseline model. An initial proposed explanation for aggressive bidding was risk aversion (e.g. Cox, Roberson, and Smith 1982). A stream of subsequent papers indicates that risk aversion parameters implied by subjects' bidding in first-price auctions do not predict behavior in other laboratory institutions. Isaac and James (2000) find that many subjects who bid as if risk-averse in a first-price auction behave as if risk-seeking in the elicitation procedure of Becker, DeGroot, and Marschak (1964). Kagel and Levin (1993) show that behavior in third-price auctions would imply risk-seeking preferences over earnings. Cason (1995) studies a first-price auction where the winning bidder pays a randomly-determined price, and finds bidding remains more aggressive than the risk-neutral prediction in that setting even though risk aversion predicts bidding less aggressively. Engelbrecht-Wiggans and Katok (2009) test the risk aversion hypothesis by providing subjects the outcomes of many independently simulated first-price auctions against robot bidders.

For completeness, we now show that maintaining constant relative risk aversion over income as a hypothesis is inconsistent with the results we report from our experiment. Therefore, one contribution of the surplus frame design is to demonstrate that risk aversion does not organize bidding behavior even staying *within* the first-price auction institution in which the winning bidder pays his actual bid when the auction is played out only once with a given realization of types.

Cox, Smith, and Walker (1988) (CSW) characterized Bayes-Nash equilibrium of the basic auction model in the case of risk-averse bidders with log-concave preferences. In particular, if all bidders have constant relative risk averse (CRRA) utility functions $u(x) = x^r$ with the same parameter $0 < r \leq 1$ for all bidders, then the equilibrium bidding function is

$$b(x) = \frac{N-1}{N-1+r}x. \tag{7}$$

The assumption of identical risk attitudes can be relaxed without changing the qualitative characteristics of the results. If bidders are assumed to be CRRA with heterogeneous parameters r_i, then each bidder's equilibrium bid function is of the form (7) for sufficiently low realizations of x_i. For large values of x_i, the equilibrium bid functions generally cannot be expressed in closed form. Van Boening, Rassenti, and Smith (1998) show numerically, under a wide range of parameter scenarios, that the most risk-averse bidders in the population have equilibrium bid functions which are concave in x_i when x_i is near the top of the support of types. These results qualitatively parallel typical behavior in the laboratory in the profit frame. Most subjects bid more aggressively than the risk-neutral prediction; most subjects bid roughly linearly in x_i over most of the support; and many subjects' bid scatterplots have a relatively flat region near the top of the support. We replicate these qualitative results in the profit frame sessions reported in this paper.

We maintain the CSW assumption that bidders are risk averse over earnings in the auction period, and derive predictions for behavior in the surplus frame. If bidders maximize the expected utility of their earnings from the transaction, with a C^2 utility function $u(\cdot)$, the maximization problem faced by bidder i is to choose her bid to solve

$$\max_{b_i} P\left(b_i\right) u\left(v - b_i\right) + (1 - P\left(b_i\right)) u\left(v - x_i\right). \tag{8}$$

The second term in the maximand in (8) captures the earnings the bidder obtains in the event she does not win the auction, and instead purchases the substitute at the price x_i.

Proposition 1. *Suppose bidder i is an expected utility maximizer over earnings in an auction period, with a C^2 utility function $u_i\left(x\right)$ satisfying $|u_i''(x)| < \infty$ for all $x > 0$. Suppose all other bidders $j \neq i$ adopt C^2 strictly monotonic bid functions b_j, with $b_j(\underline{x}) = \underline{x}$. In the surplus frame, the slope of the best-response function at type zero is independent of the utility function $u(\cdot)$. That is, bidding behavior at the low end of the support of types is approximately independent of attitudes towards risk.*

Proof. Let $P_i(b) \equiv \Pi_{j \neq i} b_j^{-1}(b)$ denote the probability bidder i will win the auction if he bids b. Given the assumptions on the bidding behavior of bidders other than i, $P(b)$ is a C^2 function with $P(0) = 0$. The first-order necessary condition for the optimal bid b at a type x is

$$P_i'(b)\left[u_i(v - b) - u_i(v - x)\right] - P_i(b) u_i'(v - b) = 0. \tag{9}$$

Let $x_i(b)$ denote the inverse best response function for bidder i. Differentiating the identity (9) with respect to b,

$$P_i''(b)\left[u_i(v-b)-u_i(v-x(b))\right]+P_i'(b)\left[x_i'(b)u_i'(v-x(b))-2u_i'(v-b)\right]+P_i(b)u_i''(v-b)=0.$$

This rearranges to

$$x_i'(b)=-\frac{P_i(b)}{P_i'(b)}\frac{u_i''(v-b)}{u_i'(v-x(b))}-\frac{P''(b)}{P'(b)}\frac{u_i(v-b)-u_i(v-x(b))}{u_i'(v-x(b))}+2\frac{u_i'(v-b)}{u_i'(v-x(b))}.$$

By substituting in (9) this can be rewritten as

$$x_i'(b)=-\frac{P_i(b)}{P_i'(b)}\frac{u_i''(v-b)}{u_i'(v-x(b))}-\left[\frac{P_i''(b)P_i(b)}{\{P_i'(b)\}^2}-2\right]\frac{u_i'(v-b)}{u_i'(v-x(b))}. \tag{10}$$

Consider taking the limit of (10) as $b\to\underline{x}$. By the assumption, $P_i(\underline{x})=0$ and $P_i'(\underline{x})>0$, and so the first term goes to zero. Because $x(\underline{x})=\underline{x}$, the limit of the second term is controlled by the quantity in square brackets, which is independent of the utility function. □

Proposition 1 applies for best responses to any fixed profile of bid functions for other bidders. Therefore, there is an immediate corollary characterizing the slope of equilibrium bid functions when bidders have identical utility functions.

Corollary 2. *Suppose all bidders are expected utility maximizers over earnings in an auction period, with the same C^2 utility function $u(\cdot)$. In the surplus frame, any symmetric equilibrium bidding strategy $b(\cdot)$ satisfies*

$$\lim_{x\to 0} b'(x)=\frac{N-1}{N}.$$

For low types, the best-response bidding behavior in the surplus frame is essentially independent of risk attitude. In particular, risk aversion cannot be cited in the surplus frame as a justification for aggressive bidding for low types. The result in Proposition 1 does not obtain for CRRA utility in the profit frame because $u'(0)$ and $u''(0)$ are not well-defined for CRRA utility. Therefore, there is a singularity at $x=\underline{x}$ in the differential equation defining the best-response bid function, and the slope at the lower boundary of the support of reservation values is not pinned down. Zero earnings is not a possible outcome in the surplus frame with $v>\overline{x}$. In fact, eliminating zero earnings as a possible outcome in the auction, in general, changes discontinuously the equilibrium bidding function for CRRA bidders. An equilibrium of the form (7) depends sensitively on zero earnings being an outcome with positive probability.

Figure 5 exhibits equilibrium bid functions for the CRRA utility functions $u(x)=\sqrt{x}$. In the profit frame, if all bidders have square-root utility over earnings, the equilibrium bid function with three bidders is linear with slope $\frac{5}{6}$. This is the median bid function slope in the profit frame sessions, and is comparable to previous findings in the literature. The surplus frame equilibrium shown is for $[\underline{x},\overline{x}]=[0,1]$ and $v=\frac{31}{30}$, which corresponds proportionately to the parameters used in the surplus frame sessions. The surplus frame equilibrium for square-root utility is slightly more aggressive than risk-neutral, but much less aggressive than the profit frame equilibrium for the same utility function. The equilibrium bidding function is convex in x, which provides a contrasting prediction to the concavity demonstrated by Van Boening, Rassenti, and Smith under the profit frame. In the surplus frame, we find that subjects bid more aggressively than the risk-neutral prediction for low types, and adopt bidding functions which are roughly linear, or concave. Both these findings falsify the predictions of risk aversion, specifically CRRA, over per-period earnings as a maintained hypothesis in organizing subject behavior.

References

[1] Armantier, Olivier, Charles A. Holt, and Charles R. Plott. 2010. A reverse auction for toxic assets. California Institute of Technology, Social Science Working Paper 1330.

[2] Becker, Gordon, Morris DeGroot, and Jacob Marschak. 1964. Measuring utility by a single-response sequential method. *Behavioral Science* 9(3): 226-232.

[3] Cason, Timothy N. 1995. An Experimental Investigation of the Seller Incentives in the EPA's Emission Trading Auction. *American Economic Review* 85(4): 905-922.

[4] Cox, James C., Bruce Roberson, and Vernon L. Smith. 1982. Theory and behavior of single object auctions. In *Research in Experimental Economics*, volume 2, edited by Vernon L. Smith. Greenwich, CT: JAI Press, 1-43.

[5] Cox, James C., Vernon L. Smith, and James M. Walker. 1988. Theory and individual behavior of first-price auctions. *Journal of Risk and Uncertainty* 1(1): 61-99.

[6] Cuzick, J. 1985. A Wilcoxon-type test for trend. *Statistics in Medicine* 4: 87-90.

[7] Engelbrecht-Wiggans, Richard. 1989. The effect of regret on optimal bidding in auctions. *Management Science* 35(6): 685-692.

[8] Engelbrecht-Wiggans, Richard and Elena Katok. 2007. Regret in auctions: Theory and evidence. *Economic Theory* 33: 81-101.

[9] Engelbrecht-Wiggans, Richard and Elena Katok. 2008. Regret and feedback information in first-price sealed-bid auctions. *Management Science* 54(4): 808-819.

[10] Engelbrecht-Wiggans, Richard and Elena Katok. 2009. A direct test of risk aversion and regret in first price sealed-bid auctions. *Decision Analysis* 6(2): 75-86.

[11] Filiz-Ozbay, Emel, and Erkut Ozbay. 2007. Auctions with anticipated regret: Theory and evidence. *American Economic Review* 97(4): 1407-1418.

[12] Harrison, Glenn W., John A. List, and Charles Towe. 2007. Naturally occurring preferences and exogenous laboratory experiments: A case study of risk aversion. *Econometrica* 75(2): 433-58.

[13] Isaac, R. Mark, and Duncan James. 2000. Just who are you calling risk averse? *Journal of Risk and Uncertainty* 20(2): 177-187.

[14] Kagel, John H. 1995. Auctions: A survey of experimental research. In *The Handbook of Experimental Economics*, edited by John H. Kagel and Alvin E. Roth. Princeton, NJ: Princeton University Press, 501-585.

[15] Kagel, John H., and Dan Levin. 1993. Independent private value auctions: Bidder behavior in first-, second-, and third- price auctions with varying number of bidders. *Economic Journal* 103: 868-79.

[16] Kirchkamp, Oliver, Eva Poen, and J. Philipp Reiss. 2009. Outside options: Another reason to choose the first-price auction. *European Economic Review* 53: 153-169.

[17] Milgrom, Paul and Robert Weber. 1982. A theory of auctions and competitive bidding. *Econometrica* 50(5): 1089–1122.

[18] Ockenfels, Axel, and Reinhard Selten. 2005. Impulse balance equilibrium and feedback in first price auctions. *Games and Economic Behavior* 51: 155-170.

[19] Roth, Alvin E. 2002. The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica* 70(4): 1341–1378.

[20] Selten, Reinhard, and Joachim Buchta. 1999. Experimental sealed bid first price auctions with directly observed bid functions. In *Games and Human Behavior: Essays in the Honor of Amnon Rapoport*, edited by D. Budescu, I. Erev, and R. Zwick. Mahway, NJ: Lawrenz Associates.

[21] Turocy, Theodore L., Elizabeth Watson, and Raymond C. Battalio. 2007. Framing the first-price auction. *Experimental Economics*, 10: 37-52.

[22] Van Boening, Mark V., Stephen J. Rassenti, and Vernon L. Smith. 1998. Numerical computation of equilibrium bid functions in a first-price auction with heterogeneous risk attitudes. *Experimental Economics*, 1(2): 147-159.

[23] Vickrey, William. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1): 8-37.

Treatment	Session	Mean winning bid	Mean relative winning bid
Profit	P1	$3.772	0.846
Profit	P2	$3.786	0.857
Profit	P3	$3.761	0.849
Surplus	S1	$3.447	0.779
Surplus	S2	$3.394	0.770
Surplus	S3	$3.356	0.757

Table 1: Summary of session-level results.

Variable	**Coefficient**	**S.E.**	p**-value**
P_{its}	0.9602189	0.0403224	
S_{its}	0.8848672	0.0112427	
tP_{its}	-0.0008072	0.0002594	0.027
tS_{its}	-0.0002686	0.0004595	0.584
$x_{its}^{MAX}P_{its}$	-0.0001896	0.0000667	0.036
$x_{its}^{MAX}S_{its}$	-0.0002414	0.0000348	0.001

Table 2: Parameter estimates for regression model (6) with errors clustered by session.

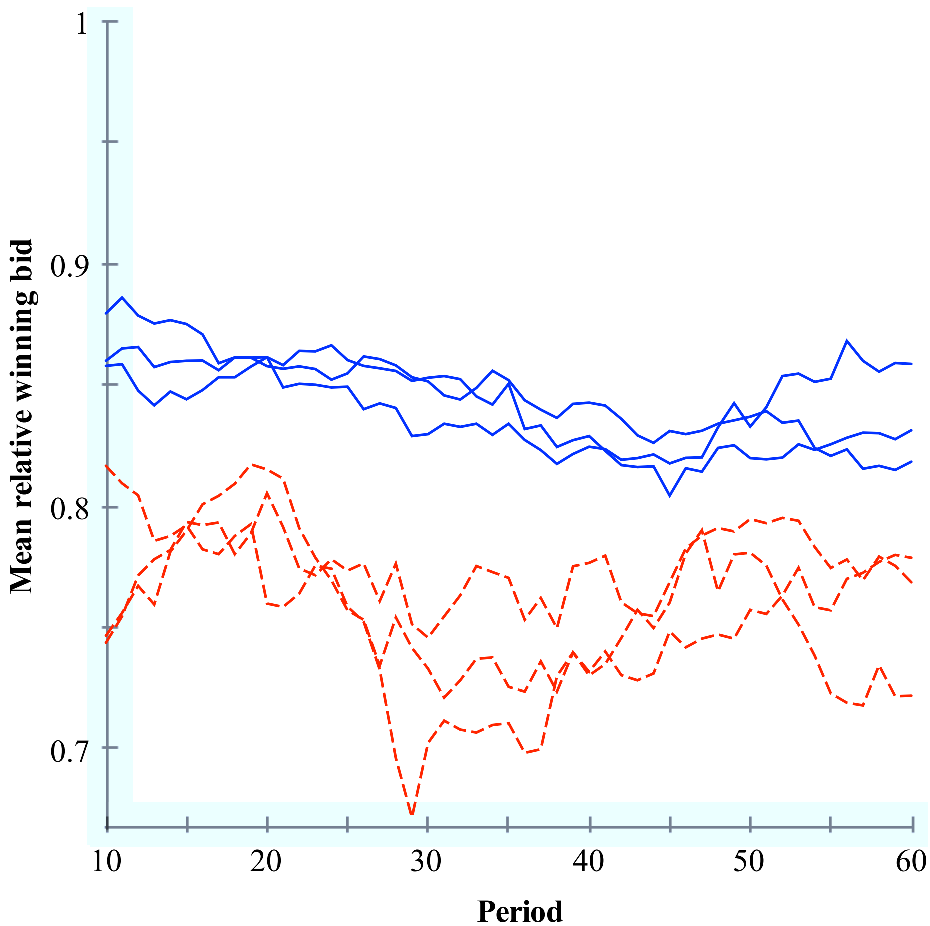

Figure 1: Ten-period moving averages of the mean relative winning bid, by session. Solid lines represent sessions using the profit frame, dotted lines those using the surplus frame.

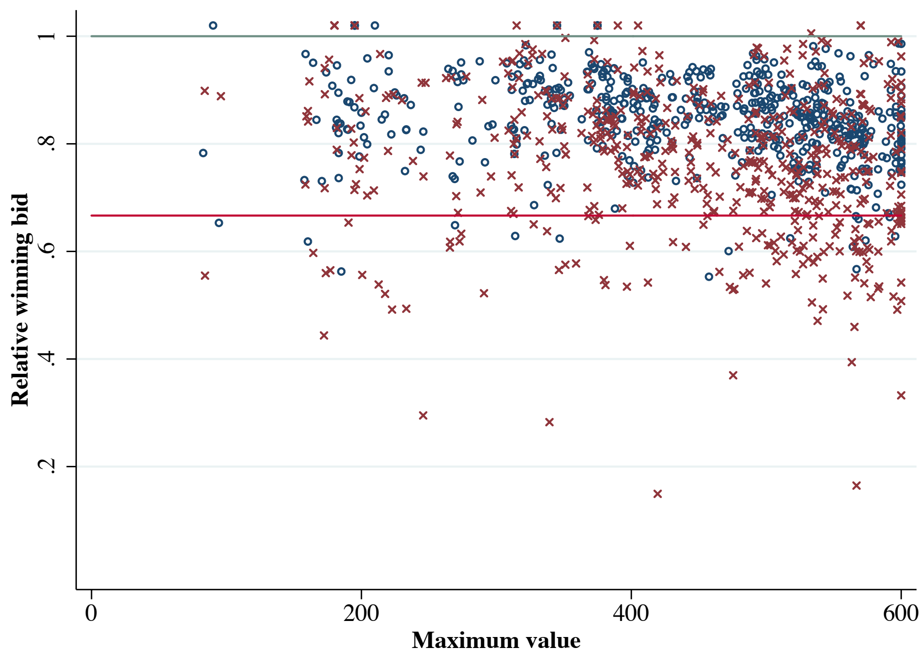

Figure 2: Scatterplot of mean relative winning bids as a function of the maximum realized value in the market, all sessions. Blue dots represent outcomes in the profit frame, red crosses outcomes in the surplus frame. Markets with winning bids above the maximum value are recorded separately in a row across the top.

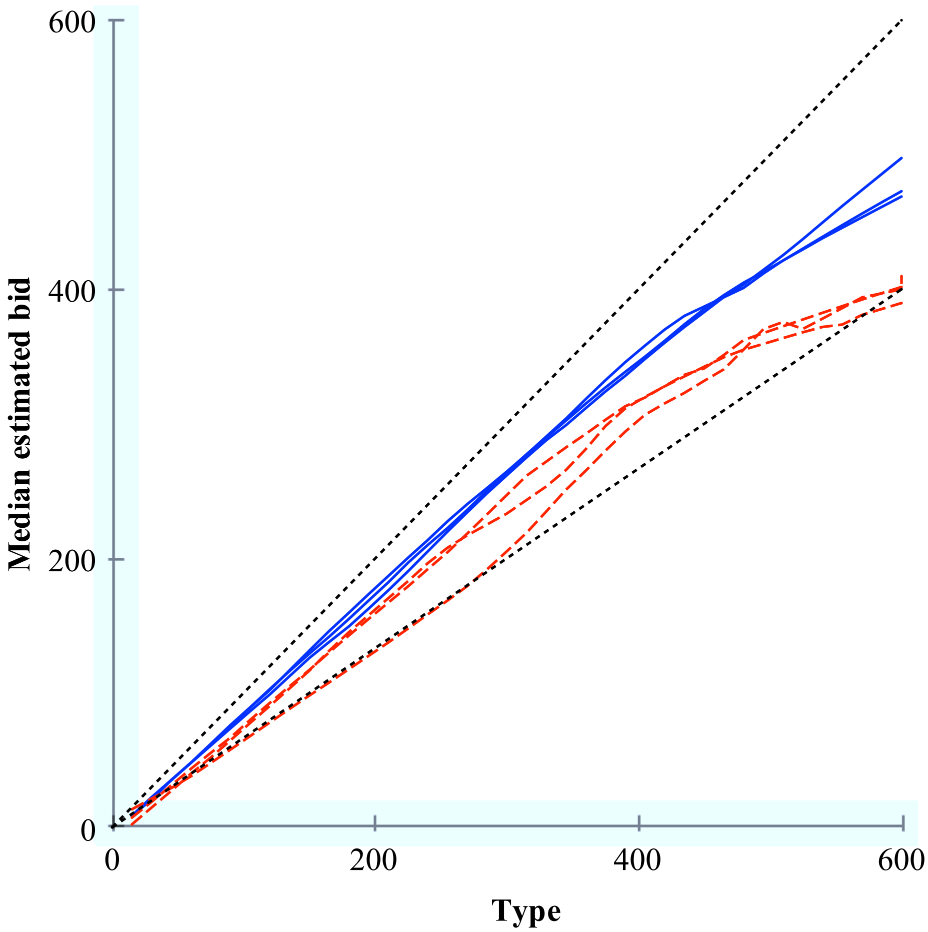

Figure 3: Median estimated bid as a function of type, for each session. Solid lines represent sessions using the profit frame, dotted lines those using the surplus frame.

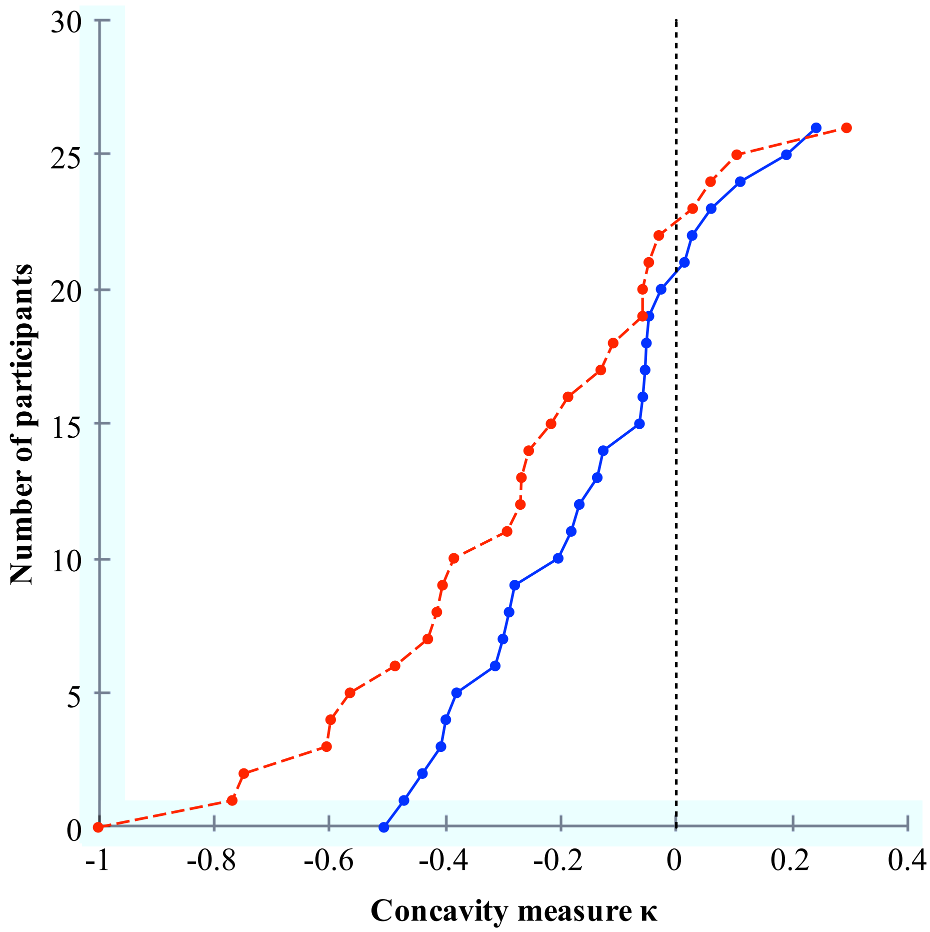

Figure 4: Cumulative distribution of concavity measure κ across bidders, by treatment. The solid line represents the distribution with the profit frame, the dashed line the surplus frame.

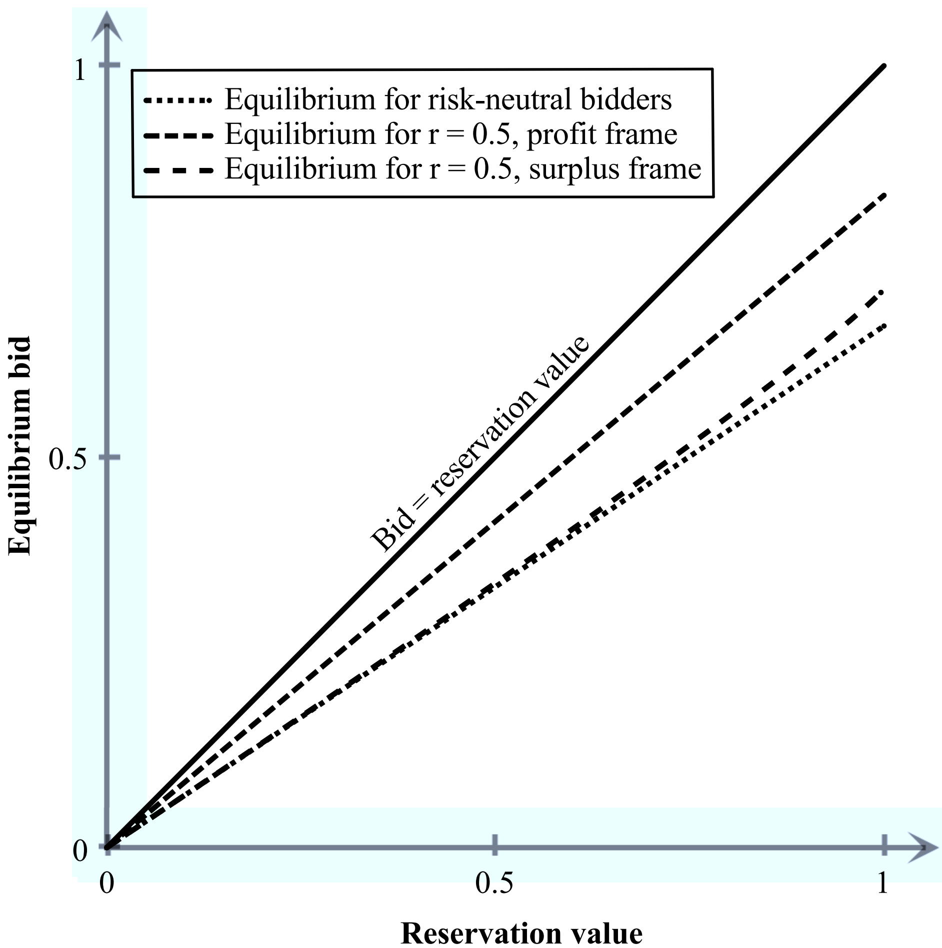

Figure 5: Comparison of symmetric equilibrium bid functions.