



City Research Online

City St George's, University of London

Citation: Georganas, S., Levin, D. & McGee, P. (2017). Optimistic irrationality and overbidding in private value auctions. *Experimental Economics*, 20(4), pp. 772-792. doi: 10.1007/s10683-017-9510-y

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/16386/>

Link to published version: <https://doi.org/10.1007/s10683-017-9510-y>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Optimistic Irrationality and Overbidding in Private Value Auctions

Sotiris Georganas

City University London

Social Sciences Building

St John Street, EC1R 0JD, London, UK

Dan Levin

Ohio State University

1945 N. High Street

Columbus, OH 43210, USA

Peter McGee

University of Arkansas

Business Building

Fayetteville, AR 72701, USA*

December 27, 2016

Abstract

Bidding one's value in a second-price, private-value auction is a weakly dominant solution (Vickrey, 1961), but repeated experimental studies find more overbidding than underbidding. We propose a model of *optimistically irrational* bidders who understand that there are possible gains and losses associated with higher bids but who may overestimate the additional probability of winning and/or underestimate the potential losses when bidding above value. These bidders may fail to discover the dominant strategy—despite the fact that the dominant strategy only requires rationality from bidders—but respond in a common sense way to out-of-equilibrium outcomes. By varying the monetary consequences of losing money in experimental auctions we observe more overbidding when the cost to losing money is low, and less overbidding when the cost is high. Our findings lend themselves to models in which less than fully rational bidders respond systematically to out-of-equilibrium incentives, and we find that our model better fits the effects of our manipulations than most of the existing models we consider.

JEL classification: C92, D44, D82

Keywords: auctions, dominant strategy, out of equilibrium payoffs

*Corresponding author, E-mail: pmcgee@walton.uark.edu, Phone:(479)575-4003, Fax:(479)575-3421. We are grateful to participants at the ESA Ft. Lauderdale Meetings, the 2009 Econometric Society Meeting, CRETE 2015, and numerous seminar participants for their helpful suggestions, and especially to John Kagel for his insightful comments. We also thank the editor and two anonymous referees for their helpful comments. Part of this research was funded by an NSF Doctoral Dissertation Research Improvement Grant.

1 Introduction

In a sealed-bid second-price auction (SPA) with private valuations, where the highest bidder wins and pays the second highest bid, bidding one's value is a weakly dominant strategy (WDS, Vickrey 1961). This strategy requires only that each bidder behave rationally and it is unaffected by the number of rivals or their valuations, a bidder's risk preferences, or beliefs regarding rationality of rivals. Repeated experimental studies have found that subjects deviate from the WDS by overbidding much more than underbidding, resulting in overbidding on average (e.g., Kagel, Harstad and Levin 1987; Kagel and Levin 1993). By contrast, experimental evidence from the strategically equivalent ascending English auction demonstrates almost immediate convergence to the dominant strategy (e.g. Kagel, Harstad, Levin 1987; Kagel and Levin 1993; Kagel 1995).¹ While overbidding relative to the risk-neutral, Nash equilibrium (RNNE) has also been frequently found in first-price auctions (FPA, e.g., Kagel 1995 and Kagel and Levin 1993), the “usual suspects,” — risk aversion, beliefs about others' play, biases in perceptions of probabilities — that may explain overbidding in FPAs are of no avail in SPAs.

The contrast between SPAs and English auctions suggests that subjects discover the WDS in the English auction but not in the SPA; why is this the case? The cognitive process that leads to the discovery of the WDS in an SPA is far from trivial and an experimental subject may be unable to recognize it without experience

¹Given the differences in the strategy spaces, it is not strictly accurate to say that the two auction formats have the same dominant strategy because bidders in an SPA are choosing a bid, where in an English auction they are deciding whether to continue or not at the current price. In the former, the dominant strategy is to bid one's value, while in the later the dominant strategy is to remain in the auction until the price surpasses one's value.

or training. In an English auction, on the other hand, a subject needs to answer a simple question for herself: “Am I ‘in’ or ‘out’?” Answering this question leads a bidder to drop out at his value.

Subjects who do not bid their value in SPAs are nevertheless still motivated by common sense economic incentives, such as expected payoffs, though imperfectly. Kagel, Harstad, and Levin (1987) conjectured that subjects are aware that higher bidding increases the probability of winning the auctions but underestimate the additional cost associated with it. Instead of looking for dominant strategies, we suggest that *optimistically irrational* bidders are guided by a desire to maximize their profits combined with an inability to fully grasp the intricacies of the auction environment that allows them to view the consequences of their actions more favorably. We do this by modeling reasonable bidders who recognize (i) a higher bid increases the probability of winning, and (ii) the bidder may understate negative payoffs to higher bids. These behaviorally plausible assumptions about bidders are the building blocks of our simple model of how out-of-equilibrium incentives might affect behavior in SPAs.

We test our model in SPAs in which we introduce a parameter that changes the expected payoff as a function of one’s bid but does not affect the WDS. The parameter multiplies realized losses by some amount β , where $\beta = 1$ is the standard case. Consistent with previous results we find that when $0 < \beta \leq 1$, overbidding is pervasive. In contrast, when we change β to 20, overbidding is significantly reduced and underbidding is more prevalent. Overbidding when $0 < \beta \leq 1$ results in very few and fairly small losses (5.8% of auctions; median loss of \$0.10). This is a product of our design: the domain of bidders’ private values is quite large relative to the

number of bidders, so the second highest bid is almost always below the highest value, even with overbidding. This allows us to rule out a “hot stove” type of learning whereby losses reduce overbidding in subsequent auctions.² Instead, it appears that the dramatic reduction in overbidding occurs when β is exogenously and publicly increased and can be attributed to changes in expected out-of-equilibrium payoffs.

While explanations for overbidding in various auction formats abound, and we compare the fit of our model to several of them in Section 5, the contribution of our model lies in its focus on the dominant strategy, adding to the recent theoretical interest in how dominant strategies influence decision calculus in games (e.g., *obvious strategy-proofness*, Li 2016). Our strong findings suggest that incentives outside equilibrium affect behavior in predictable ways in the laboratory, and probably in the field as well, even when equilibrium analysis predicts otherwise. Goeree et al. (2002) show a similar result in a FPA, but in FPA, as in many other games where Nash equilibrium is the solution concept, best responding requires “cardinal” computations. Since such computations often involve a high degree of complexity and a heavy mental cost, we do not expect that the outcome in FPA will exactly reflect

²The change in behavior we observe is immediate, once the β parameter changes, and more extreme than could be plausibly predicted by learning models, as they would usually be applied (note that little is known about learning transfer between different but very similar games, such as the ones that result from the manipulation of β). In particular, reinforcement learning (Erev and Roth 1998) would predict no change in behavior before subjects have a chance to experience the new payoffs, unlike what we observe in the data. Fictitious play (Fudenberg and Levine 1998) is also unlikely to fit the data, since the overbidding we observe can not be justified by any beliefs. Other models, such as learning direction theory (Selten and Stoecker 1986) and EWA (Camerer and Ho 1999), consider foregone payoffs. These might predict a faster response in auctions than other models, but it is still hard to conceive how they could predict such an immediate change in bidding, when β changes, after the subjects have already been learning for 20 or 40 periods. Steady state concepts such as QRE (McKelvey and Palfrey 1995) are better suited to the task, e.g. predicting the immediate tendency towards value bidding when β rises without a need for subjects to actually experience losses. We calculate and fit such models in Section 5.

the point prediction of Nash equilibrium. It is much less surprising to find that the subjects’ calculations, possibly involving heuristics, approximations and simplification rules, will be affected by a change in the incentives, even if these ought to have no effect on Nash equilibrium. This complexity motivates many models that predict overbidding by allowing bidders to make, and learn from, mistakes (e.g., QRE). In an SPA with private values, however, the dominant strategy can be reached with just “ordinal logic” of dominance, without even a need for common knowledge of rationality.³ Thus, one would expect the solution norm— bid your value —to have its best chance for success in this environment. Our study shows that behavior is still guided by some degree of conscious profit maximization, but subjects’ decision processes fail to recognize a characteristic that is very seldom present outside the lab: the dominant strategy. Errors in recognizing a dominant strategy require a new perspective on the cognitive processes underlying bidding behavior of the sort provided by *optimistic irrationality* to try to explain “errors” made by bidders that are as much a function of the simplicity of a dominant strategy as the complexity of the environment.

2 Optimistically Irrational Bidders

The overlooked availability of a WDS must be the starting point of any explanation to overbidding in SPA. We formalize the intuition behind the conjecture laid out in

³Levin, Peck and Ivanov (2016) is devoted to separating failures of “ordinal logic” (or “insight”) from the failures of and/or biases in Bayesian updating. They use tasks involving colored cards to measure these failures and biases, and use the measures to analyze deviations from predicted bidding behavior in a common-value Dutch auction.

Kagel, Harstad, and Levin (1987) by modeling an “*optimistically irrational*” bidder who understands that there are possible gains and losses associated with higher bids but who may overstate the additional probability of winning due to higher bidding and/or understate the losses associated with it.⁴

Let there be n risk-neutral bidders, each of whom privately observes her value x_i , $i = 1, \dots, n$. It is common knowledge that the x_i 's are *i.i.d* draws from a distribution with a cumulative density function $F(t)$, where $F'(t) = f(t) > 0$ on $[0, 1]$, $F(0) = 0$, and $F(1) = 1$. For our purposes, we assume that the x_i 's are drawn from a *generalized uniform* distribution, $F(t) = t^{\hat{\alpha}}$ and $\hat{\alpha} \geq 1$, where $\hat{\alpha} = 1$ corresponds to the *uniform* distribution used in almost all laboratory SPAs and FPAs. We first make four assumptions:

Assumption 1. *Symmetry*

Assumption 2. *Upon winning, a bidder's gross payoff is x .*

Assumption 3. *Bidders believe values are i.i.d from a c.d.f. $F(t) = t^\alpha$ with $\alpha \geq \hat{\alpha}$.*

Assumption 4. *A bidder with a value x who bids $b > x$ and wins at a price $p \in (x, b)$ believes that the expected payment is $\gamma(x)\frac{b+x}{2}$, with $0 < \gamma(x) \leq 1$, i.e., expected losses are $x - \gamma(x)\frac{b+x}{2}$.*

The first two assumptions simply mean that each bidder believes that all other bidders use the same, strictly monotonic bidding function, and that they receive their

⁴The conjecture is as follows, “Bidding in excess of x in the second-price auctions would have to be labeled as a clear mistake, since bidding x is a dominant strategy irrespective of risk attitudes. Bidding in excess of x is likely based on the illusion that it improves the probability of winning with no real cost to the bidder as the second-high-bid price is paid.” (Kagel, Harstad, and Levin, p. 1299, 1-5)

full value if they win the auction. Assumption 3 implies that a bidder potentially overstates the impact of bidding past his value because he believes values are closer together than they actually are. That is, he believes the increase in the probability of winning corresponding to an increase in his bid is greater than it actually is. This assumption finds support in other studies of auctions. For example, Cooper and Fang (2008) found that subjects who perceive their rivals to have values similar to their own are more likely to overbid in experimental SPAs, while Breitmoser (2015) uses “projection,” the tendency to believe that rivals have value or beliefs similar to one’s own, to explain the winner’s curse.⁵ Finally, Assumption 4 captures the notion that a bidder in an SPA may understate possible losses when he bids above his value *and* wins.

These assumptions allow us to make two important observations about an optimistically irrational bidder’s maximization problem

Proposition 1. *Sincere bidding, i.e., $b(x) = x$ for all x , occurs if and only if for all x , $\gamma(x) = 1$.*

Proof. See appendix. □

Proposition 1 means that an optimistically irrational bidder may still bid his value when $\alpha = \hat{\alpha}$ and $\gamma(x) = 1$. This does not require that the bidder recognizes the availability of the WDS. In fact, sincere bidding only requires $\gamma(x) = 1$ but allows $\alpha > \hat{\alpha}$.⁶

⁵Perception of similarity in Cooper and Fang is induced by information provided to subjects about their rivals’ values.

⁶When $b(x) > x$, the value of α will, however, affect the extent one bids above his value.

Remark. *There is a linear solution to the maximization problem with $b(x; \alpha, \gamma, n) = \delta(\alpha, \gamma, n)x$.*⁷

Put more simply, the bid will be a multiple of the value. The exact multiple will be a function of the number of competitors, how “close” a bidder believes those competitors’ values are to his value, and the extent to which he understates the losses from overbidding. In the next section we make use of the fact that there is a linear solution to the maximization problem to inform our experimental design.

2.1 Experimental Test of the Model

Optimistically irrational bidders need not recognize the WDS but can be influenced by “out-of-equilibrium” payoffs. Several authors have argued that, in English or first-price auctions with private values, the exact shape of the expected payoff functions matters (e.g., Harrison 1989; Goeree et al. 2002; Georganas 2011; Georganas and Nagel 2011).⁸ Similarly, Noussair, Robin, and Ruffieux (2004) compare value revelation using the Becker-DeGroot-Marschak (BDM) method and SPAs and find that the shape of the expected payoff function may influence behavior. They find that the shape of the expected payoff function in the SPAs means that the probability of winning increases faster in overbids than it decreases in underbidding, which drives bidders who start out below the WDS towards the WDS, while overbidding is also more costly in the SPA than in the BDM with three or more bidders. Taken together, these observations about the shape of the expected payoff function seem

⁷The formal derivation can be found in the appendix.

⁸Our model of optimistically irrational bidders also predicts bidding above the risk-neutral Nash equilibrium in FPAs that is affected by the shape of the expected payoff function, which we return to in Section 5. Details can be found in the appendix.

to reflect the assumptions we make in our model. As a simple test of optimistically irrational bidders, we implement standard second-price, sealed-bid auctions for one unit of an indivisible good in the laboratory. The WDS predicts that players bid their values in equilibrium. We introduce a factor β by which we multiply eventual negative profits of the bidders. β does not affect the equilibrium if bidders are bidding sincerely because no bidder earns negative profits in equilibrium. However, the bidders' expected payoff functions do change, given that their opponents follow the equilibrium strategies; in Figure 1 we plot, for different values of β , the expected payoff function for a bidder whose rivals bid their values.

Given the remark in the previous section, we can determine how our experimental treatments will affect the bidding of optimistically irrational bidders. Applying the theory to our experimental treatments, when a bidder wins with $b > x$ and loses, his expected payment is $\beta[x - \gamma_\beta \frac{b+x}{2}]$. A standard SPA corresponds to the situation where $\beta = 1$, such that the expected losses are $1 \times [x - \gamma_1 \frac{b+x}{2}]$. This allows us to define γ_β as a function of γ in a standard auction, i.e., γ_1 , which we do by setting $[x - \gamma_1 \frac{b+x}{2}] = \beta[x - \gamma_\beta \frac{b+x}{2}]$, and solving for γ_β . Doing this, we get

$$\gamma_\beta = \frac{\beta - 1}{\beta} \frac{2x}{b+x} + \frac{\gamma_1}{\beta} \quad (1)$$

We can substitute $b(x) = \delta x$ into equation (1) to obtain

$$\gamma_\beta = \frac{1}{\beta} \left[\frac{2(\beta - 1)}{\delta + 1} + \gamma_1 \right] \quad (2)$$

Plugging equation (2) into the linear solution, we solve the problem numerically.⁹ The results can be found in Table 1. Assuming $\alpha = 1.1$ and $\gamma_1 = 0.95$, i.e., a bidder slightly overstates the increase in his probability of winning associated with a higher bid and slightly understates his expected loss, the model predicts that δ^* increases monotonically in β and bidding above value that ranges from 9.7% above value for a $\beta = 0.1$ to 0.5% above value for a $\beta = 20$. These numerical estimates yield three testable hypotheses:

Hypothesis 1. *Subjects will overbid on average*

Hypothesis 2. *Subjects will overbid by more on average when $\beta < 1$ than when $\beta = 1$*

Hypothesis 3. *Subjects will overbid by less on average when $\beta > 1$ than when $\beta = 1$*

3 Experimental Details

The data come from nine experimental sessions conducted at Ohio State University. Students were recruited via e-mail and sessions took place in the Experimental Economics Lab. The experiment was programmed and conducted with the software z-Tree (Fischbacher 2007). In every session, subjects participated in 62 second-price, sealed-bid auctions—2 trial auctions followed by 60 paying auctions—with

⁹The linear solution can be found in the appendix.

either three or six bidders per auction. Subjects were randomly and anonymously re-matched between auction periods.¹⁰ At the beginning of each auction, subjects privately observed their own independent private values denominated in an experimental currency unit (ECU), but they did not observe the values of others. All values were drawn from a uniform distribution on the interval $[0, 100]$, which was common knowledge. At the end of each period the bidder who obtained the item was informed of the price and his profit, while bidders who did not obtain the item received no information about the price or the bids of others. The instructions can be found in the appendix.

We multiplied negative profits by a parameter, β , which took on three values in every session: 1, 0.1, and 20. Beta took on one of these values for periods 1-19, another for periods 20-39, and the final value for periods 40-60. Subjects knew the value of β , that it is the same for all bidders, and that they would be made aware when it changed; subjects did not know when β would change, how many times it would change, or what its magnitude would be. All subjects were given starting balances of 150 ECUs to cover the possibility of losses. Profits and losses were added to this balance and the balance was paid at the end of each session. We ran two sessions with 3 bidder auctions and a β order of 1, 0.1, 20 ($\beta_{1/0.1/20}^3$), two sessions with 6 bidder auctions and a β order of 1, 0.1, 20 ($\beta_{1/0.1/20}^6$), two sessions with 3 bidder auctions and a β order of 1, 20, 0.1 ($\beta_{1/20/0.1}^3$), and three sessions with 3 bidder auctions and a β order of 20, 1, 0.1 ($\beta_{20/1/0.1}^3$). In sessions with three bidder auctions, the exchange rate was \$1=20 ECU, while the exchange rate in sessions

¹⁰Sessions had between 12 and 24 subjects each.

with six bidder auctions was \$1=14 ECU. The exchange rates were different in order to equalize the expected payoff between sessions with different group sizes. In the event that a player went bankrupt, they were no longer permitted to bid and were paid a participation fee of \$8. Due to the uneven numbers after a bankruptcy, at the beginning of every period after a bankruptcy two subjects were randomly assigned to sit out that period in 3 bidder auctions, while five subjects were chosen to sit out in 6 bidder auctions. In the six sessions that started with $\beta = 1$, there were two bankruptcies; in the three sessions that started with $\beta = 20$ there were 10 bankruptcies, with all but one occurring in the first 10 periods. Complete session details can be found in Table 2.

4 Results

Average differences between values and bids can be found in Table 3, and Figures 2-6 compare subjects' bids and their values. Consistent previous research, we see overbidding on average in every treatment for every value of β . To test whether or not this overbidding is significantly different from 0, we calculated the mean difference between bid and private value for each session within a β regime, i.e., the block of periods during which β remained the same. Using these means as our measure of overbidding, average overbidding is significantly greater than 0 at the 5% level in every case except for $\beta = 20$ for $\beta_{1/0.1/20}^3$ (t-test $p = 0.174$) and $\beta_{1/0.1/20}^6$ (t-test $p = 0.129$).

Observation 1. *Consistent with Hypothesis 1, subjects overbid on average for all*

values of β .

The effects of changes in β are visible and drastic, in contrast to the standard theoretical prediction of no change at all. For example, when we reduce the punishment for negative outcomes from $\beta = 1$ to $\beta = 0.1$ in period 20 of $\beta_{1/0.1/20}^3$, there is an immediate effect as the average difference between bid and value more than doubles from approximately 2.8 to 7.5— an increase equal to approximately 5% of the support of values. When β rises to 20 in period 40 and punishment for negative outcomes is severe, the overbidding largely disappears, with average overbids falling from 7.5 to 1.1. Similar patterns emerge in all treatments, and the differences in overbidding across β regimes are significant.^{11,12}

Observation 2. *Consistent with Hypotheses 2 and 3, overbidding varies significantly across different values of β , with higher levels of β leading to less overbidding.*

Our specific design allows for an additional important and interesting observation about learning. There is significant overbidding in the first 40 periods with $\beta = 1$ and $\beta = 0.1$ in $\beta_{1/0.1/20}^3$ and $\beta_{1/0.1/20}^6$, and we do not observe learning in the direction of value bidding in these periods.¹³ Nonetheless, there is a drastic reduction of

¹¹To compare overbidding, we conducted Wilcoxon signed rank tests on the session means we calculated for the various β regimes. In pairwise tests of each value of β for all treatments, all differences were significant at the 5% level or better except for $\beta = 1$ versus $\beta = 0.1$ for $\beta_{1/20/0.1}^3$ ($p = 0.454$).

¹²Merlob, Plott, and Zhang (2012) compare auction designs used for Medicare and Medicaid procurement auctions. In auction designs that allow for costless renegeing, which is similar to bidding in an auction with $\beta < 1$, they also find overbidding.

¹³We divide those periods for which $\beta = 0.1$ in these treatments, i.e., periods 20-39, into 4 blocks of 5 periods each. We then calculated session specific mean overbidding for each block. The means of these means for the blocks are 2.18, 2.40, 2.83, and 4.60, respectively. Although the means are increasing, pairwise Wilcoxon signed rank tests of these means reveal no significant differences between any of the first three blocks, while the last block is significantly different from all the others but further from the WDS.

overbidding in period 40. One possible explanation for the decline in overbidding in period 40 is that subjects who have overbid in earlier periods are chastened by losing money, a sort of “hot stove” learning. The evidence in Table 4 does not support this explanation. Before β increases to 20, few auctions result in realized losses. They range from a minimum of 4.3% of auctions for the case with six bidders when $\beta = 1$, to 9.3% with three bidders and $\beta = 0.1$. The average loss is also quite small, ranging from 1.1 ECUs with six bidders and $\beta = 0.1$, to 10.8 ECUs with six bidders and $\beta = 1$. Moreover, this small and infrequent negative feedback for the first two levels of β appears to have no effect on those who experience it. Of the 31 bidders who lose money when $\beta = 1$, 29 also lose money when $\beta = 0.1$. Of the six bidders who lost money when $\beta = 20$, 5 lost money at all three levels of β and the sixth lost money when $\beta = 0.1$.

Observation 3. *Reductions in overbidding when β is increased are not caused by learning due to previous losses.*

In order to move beyond unconditional means and investigate the effects of β while allowing for individual heterogeneity, we estimate a random effects Tobit model, regressing the difference between subject i 's bid and his value in auction j on dummies for each β regime.¹⁴

The results in columns 1-4 of Table 5 are similar to the means in Table 3. In every treatment bids are significantly higher when $\beta = 0.1$ than when $\beta = 1$; similarly, bidding is significantly lower in every treatment when $\beta = 20$ than when $\beta = 1$. We do, however, see some differences across treatments, and we reject the null hypothesis

¹⁴In the experiment, subjects could not bid more than 100, which can be seen in the groups of bids clustered at 100 in Figures 3 and 4. We use a Tobit model to account for the censoring of bids.

that the marginal effects of $D_{\beta=0.1}$ and $D_{\beta=20}$ are jointly equal across all treatments (Wald test, $p = 0.000$ for both).

Although learning due to negative reinforcement is unlikely in $\beta_{1/0.1/20}^3$ and $\beta_{1/0.1/20}^6$, one possible explanation for the differences across treatments we observe may be some other sort of learning or experience. To address this possibility, we augment the model in columns 5-8 with a linear time trend and its interaction with the β regime. We observe small, but statistically significant increases in bids over time in the first 20 periods in treatments $\beta_{1/0.1/20}^6$, $\beta_{1/20/0.1}^3$, and $\beta_{20/1/0.1}^3$, where bids increase by roughly 0.12 to 0.18 ECUs per period; we see no significant changes over time in subsequent periods. With the inclusion of the time trend, the same pattern emerges: greater overbidding when $\beta = 0.1$ and less overbidding when $\beta = 20$, though the treatment effects are no longer significant in $\beta_{1/20/0.1}^3$. Allowing for changes to bidding behavior over time, we cannot reject the null hypothesis that the marginal effects of $D_{\beta=0.1}$ and $D_{\beta=20}$ are the same across all treatments (Wald tests, $p = 0.551$ and $p = 0.318$, respectively); we find no significant differences at conventional levels across treatments in pair-wise comparisons of $D_{\beta=0.1}$ and $D_{\beta=20}$.^{15,16}

Observation 4. *The effects of changes in β are robust to controls for individual differences and learning over time.*

¹⁵The only pairwise comparisons with p-values less than 0.2 are comparisons of $D_{\beta=20}$ in $\beta_{1/20/0.1}^3$ with all other treatments.

¹⁶One possible motivation for overbidding not found in Table 4 is spite. Cooper and Fang (2008) find that overbidding can be attributed to both a heightened sense of competition for those with high values, and spiteful behavior by those with low values. We see little evidence of spite. In results available from the authors, we ran the same models as in columns 5-8 but restricted observations to only those with values in the top 75% or the bottom 25% of the support of values. The results for either sample are similar to those in table 5. Moreover, among bids between 95 and 100, only 12 of 867 bids come from bidders with values in the bottom quarter of the support.

The linear trend presupposes that the effect of all periods is the same within a β regime, however Figure 2 suggests that the first few periods in a session might be slightly different.¹⁷ In columns 9-12 we estimate the same model as in columns 5-8 but exclude the first 3 periods in each session. After excluding the first three periods, there is no significant learning over time for any value of β in any treatment.

Observation 5. *Bidding behavior evolves substantially in the first few periods of a session, but little thereafter.*

5 Alternative models

One important test of our model of optimistically irrational bidders is how well it fits the data relative to existing models. Among our candidate models, we begin by considering a symmetric Nash equilibrium (SNE) with normally distributed errors, given that without errors the SNE fails completely to predict the change in bidding when we shift β . Models that take all payoffs into account, even if they are not on the equilibrium path, are good candidates to explain our results. Perhaps the simplest way to consider all payoffs is to use the Nash model but assume that subjects' errors depend on the expected utility of each action in a systematic way. We do this by considering an SNE with a logistic error structure, as in Crawford and Iriberri (2007). Finally, we consider is quantal response equilibrium (QRE), which also makes explicit use of the payoff function shapes, by positing that players choose an action with a

¹⁷All sessions began with either $\beta = 1$ or $\beta = 20$, so the relevant data are the first few data points in the first and third panels of Figure 2. Similar to Noussair, Robin, and Ruffieux (2004), we observe bidding below the WDS at first before bids rapidly increase.

probability proportional to its expected payoff.¹⁸ In preliminary comparisons, we find that the SNE+normal model outperforms the other two models in all but one β -number of bidder combinations.¹⁹ The reason is that, under a logit error structure, a high frequency of underbidding is predicted for intermediate private values, since the expected payoffs are quite flat to the left of the maximum (as seen in Figure 1); yet we see far more overbidding than underbidding in all cases.²⁰ QRE improves on the Nash model with logistic errors, but still performs worse than the Nash model with normal errors.

One way to account for the fact that we observe much more overbidding than underbidding is to allow bidders to experience joy-of-winning (JOW). JOW can be incorporated by adding an extra fixed utility, U_i , to the payoff of subject i , conditional on winning the auction.²¹ It is easy to show that with such modification a

¹⁸Another possible model that we do not consider is a level-k model of the sort used to explain auction results in Crawford and Iriberri (2007). This model predicts neither overbidding nor reaction to different values of β , even allowing symmetric errors. In such a model, players of level one or higher are characterized by different beliefs, but they all best respond to these beliefs. Since bidding one's value is a weakly dominant strategy in a SPA and does not depend on a player's beliefs (or risk attitudes), players of all levels are predicted to bid their values, as in any Nash equilibrium.

¹⁹Table A1 in the appendix provides the details of the comparison of fits among all alternative models we consider.

²⁰For example, for a value of 50 and $\beta = 0.1$ we should see approximately the same amount of underbidding as overbidding, which is clearly not the case.

²¹Note that JOW looks superficially similar to spiteful behavior (e.g., Andreoni et al. 2007, Cooper and Fang 2008). The difference is that a subject experiences JOW only in the case where she wins, while a subject exhibits spite even if she does not win but in cases where she raises the price for other bidders. In Cooper and Fang's experiment, for example, subjects can sometimes acquire information about their rival's value. They find that, when subjects know that their value is close to their rival's, overbidding is more likely to result in a costly loss and subjects overbid by less. Conversely, when subjects know that their value is well below their rival's, hence overbidding is unlikely to result in winning the auction but can raise the price for their rival, overbids are larger. This is not dissimilar to our manipulation of β , and Cooper and Fang suggest that this behavior is consistent with a modified JOW/spite model. Even the modified model cannot help explain the difference between SPAs and English auctions, and Figures 3, 4, and 5 reveal that the constant overbidding predicted by JOW is more plausible than spite in our data, so we consider only JOW.

new dominant solution emerges, with $b_i(x_i) = \frac{U_i}{\beta} + x_i$.²² This implication helps as it predicts that players who enjoy winning will overbid with respect to the Nash equilibrium and the amount of overbidding will depend inversely on β . The JOW parameter j is found to be positive and yields a significantly higher likelihood in every case. Nonetheless, SNE with normal errors still provides the best fit among JOW models.

To evaluate its broad applicability, in Table 6 we examine how well the SNE models with normal errors, both with and without JOW, fare against our model of optimistically irrational (MI) bidders in SPAs, English auctions, and FPAs, by comparing estimated log-likelihoods.^{23,24} We find that our model fares better than the SNE with just normal errors (but no JOW) in all auctions. On the other hand, the SNE with normal errors and JOW outperforms our model in every auction. The SNE with normal errors and JOW may have slight advantages by the measure of Table 6, but these advantages do not reveal the full story. In Table 7 we break out the fit by the number of bidders and values of β . In this case, the SNE fares better than MI in six of the nine β -number of bidder combinations. In Table 8, we compare the predicted mean overbidding by optimistically irrational bidders and the SNE with JOW to the observed mean overbidding. In 4 out of the 7 cases, the magnitude of overbids by MI

²²Details are provided in the appendix.

²³The English auction data come from Georganas (2011), and the FPA data come from Dyer et al. (1989).

²⁴The derivation of optimistic irrationality for FPAs can be found in the appendix. The implications of optimistic irrationality for bidding in English auctions are much more straightforward. Assumption 4 of the model lays out the bidder's beliefs about expected payments and losses if he wins at a price greater than his value, but there is no uncertainty over payments or losses in the English auction, so beliefs should be perfectly accurate, leading even an optimistically irrational bidder to bid his value.

bidders is closer to the predicted overbidding than SNE bidders who experience JOW. While on a strictly econometric basis, SNE+JOW seems to be performing slightly better than MI, MI outperforms SNE+JOW in several instances.²⁵ Moreover, there are good qualitative reasons not to be satisfied with JOW, chief among them its failure to explain why JOW occurs in SPAs but not in the strategically analogous English auction. Ultimately, further work and additional data will be needed to completely analyze the relative strengths of the two models.

6 Conclusions

Experiments consistently find that in second-price, sealed-bid auctions with private values— a mechanism with incomplete information where bidding one’s value is a WDS — subjects deviate significantly from the WDS. The availability of a “dominant” action that is best irrespective of the other features of the decision is rare in games with incomplete information and in strategic situations outside the lab. The behavior of a bidder in a second-price auction who fails to recognize or discover such an available strategy is still likely to be guided by rules that are useful in a wide range of situations, such as cost-benefit analysis. Subjects in our SPAs provide support for this characterization of bidders: their bidding is reasonable if not optimal. Subjects overbid on average but their overbidding is influenced by manipulations which affect expected payoffs out of equilibrium but not the dominant strategy. In accordance with lessons learned in more familiar settings, as we vary the magnitude of

²⁵Although we do not include the SNE model with normal errors and JOW in Tables 7 and 8, MI outperforms that model in every case.

the penalty for losses, a natural reaction is to hedge and bid lower when the penalty is relatively larger and to be more aggressive when the penalty is relatively lower. The behavioral changes may not be optimal in a second-price auction, yet they are sensible when viewed through the lens of their applicability in richer environments.

We propose a model of *optimistically irrational* bidders who fail to recognize the availability of a dominant strategy. Bidders in this model understand that raising their bid increases the probability of winning but may either overstate the increase in the likelihood of winning and/or fail to appreciate the costs associated with increasing their bids. We fit several existing models designed to explain overbidding using our data, but we find that most of these models perform poorly even when they consider out-of-equilibrium payoffs that would be affected by our experimental manipulation, and none of the models outperform ours consistently.

Our results build on the cautiously optimistic findings in Cooper and Fang (2008). They find that bounded rationality—more than non-standard preferences like spite and JOW—contributes to overbidding in SPAs. Subjects in their experiment could purchase costly and noisy information about rivals' values, information which does not affect the WDS. Subjects who purchase the information were significantly more likely to overbid, but the behavior of those subjects who did not purchase information was consistent with theoretical predictions. They conclude by noting that this heterogeneity may be of less significance outside the lab where selection might weed out the irrational bidders, leaving only rational bidders. Our finding of large and theoretically unpredicted responses to our treatments can inform mechanism designers, theorists and practitioners who are concerned that such selection may be insuffi-

cient or too slow: even in cases with a dominant strategy, the nature of incentives outside equilibrium can influence behavior. In instances where the common sense implications of manipulating out-of-equilibrium incentives can steer behavior toward the desired norm, such as in SPAs, designers may be able to use these incentives to design more stable and efficient mechanisms.²⁶

²⁶While our results are most applicable to situations with a dominant strategy, the fundamental intuition behind optimistically irrational bidders may be useful in other domains, as well. For example, over investment is common in experimental contests (Dechenaux, Kovenock, and Sheremeta 2015) that test models of RD competitions. Optimistically irrational competitors may overestimate the increase in the likelihood of winning an RD contest and/or underestimate just how costly their effort may be.




Figure 1: Expected payoff functions in a second price auction with 3 bidders, for the three different values of β . There are 5 curves in every panel which represent expected utility, depending on one's bid for a private value $v = 0, 25, 50, 75$ and 100 .

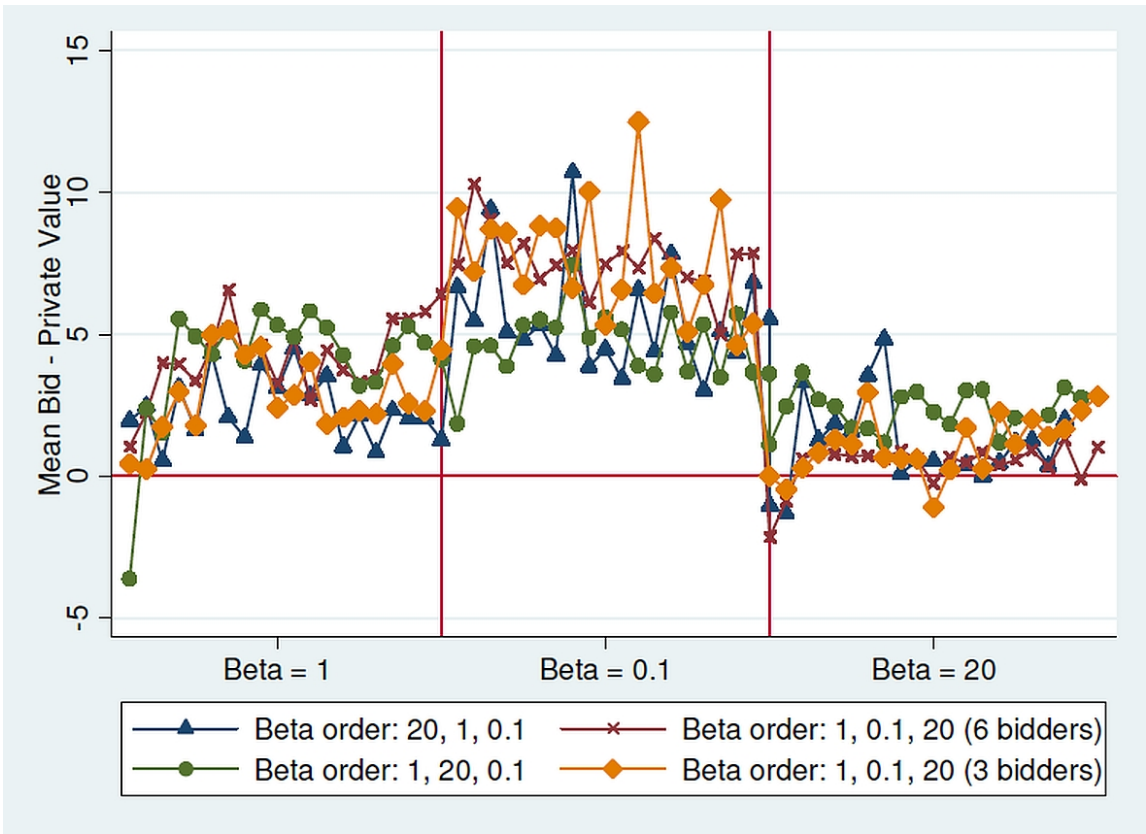


Figure 2: Mean difference between bids and private values over time in all treatments

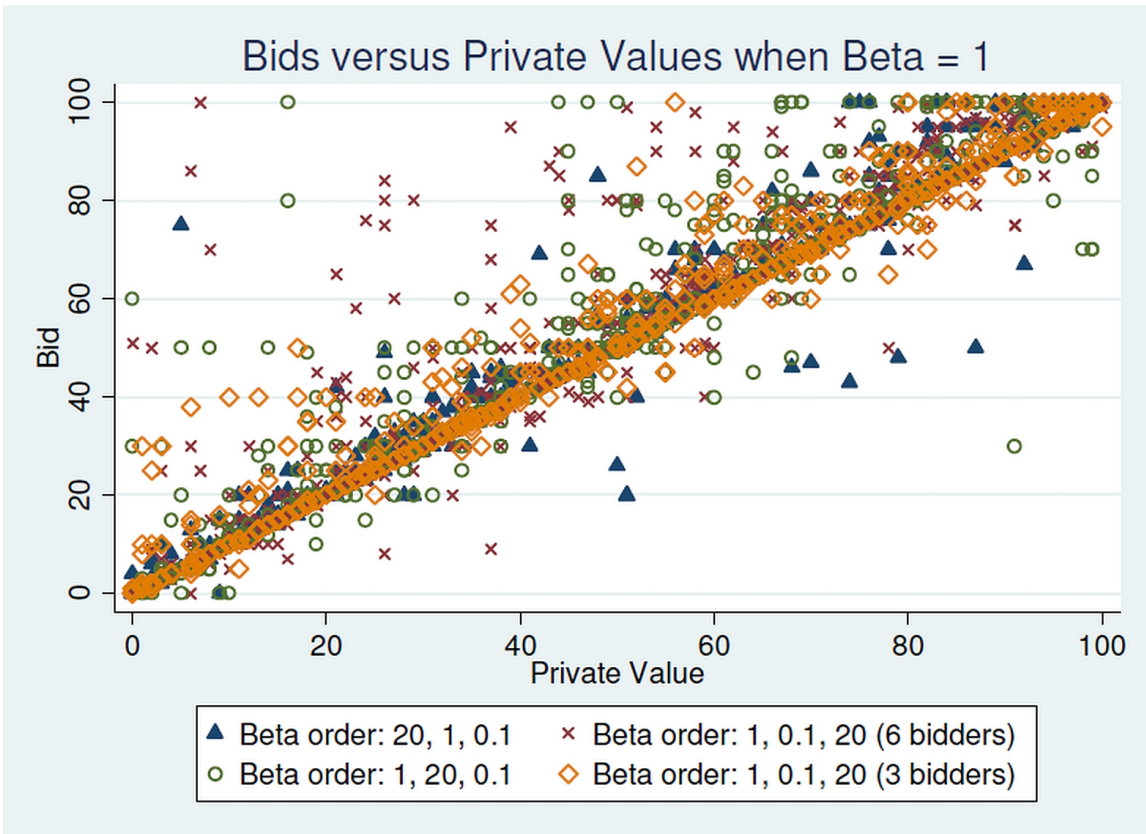


Figure 3: Scatter plot of values versus bids in all treatments when $\beta = 1$

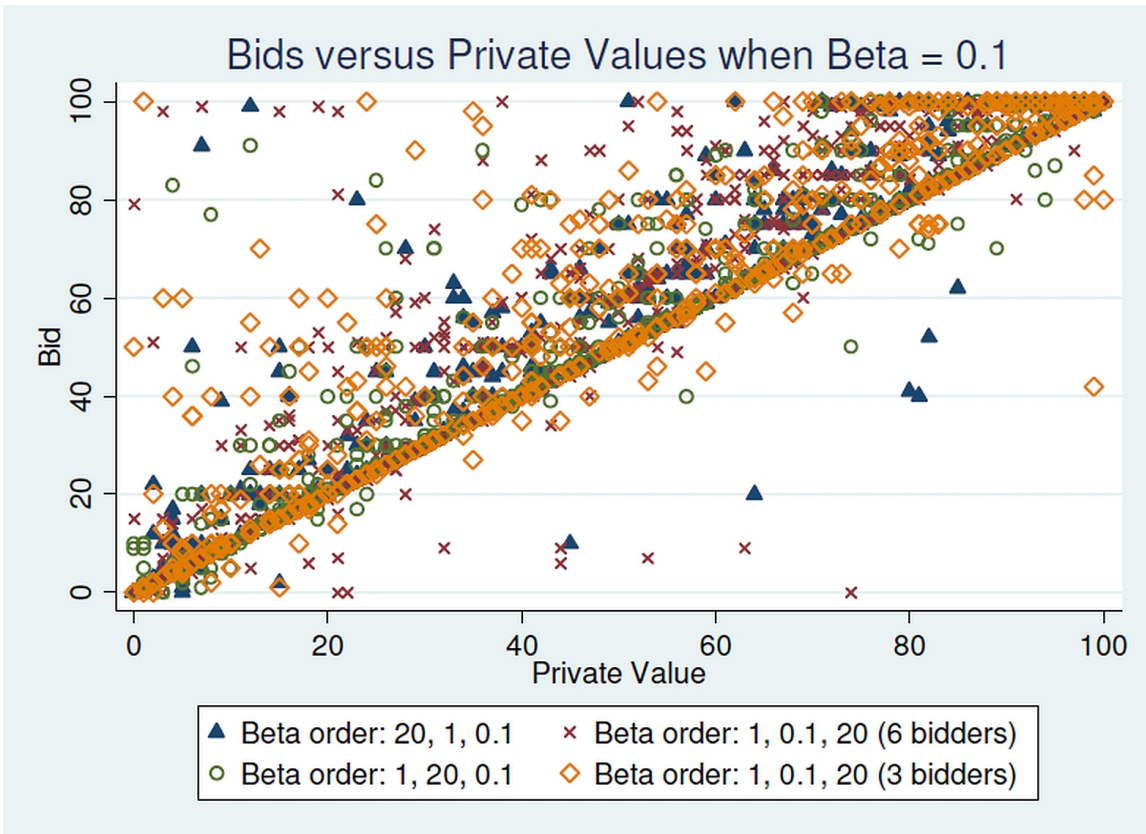


Figure 4: Scatter plot of values versus bids in all treatments when $\beta = 0.1$




Figure 5: Scatter plot of values versus bids in all treatments when $\beta = 20$

Table 1: Numerical estimates of parameters for optimistically irrational bidders in auctions with 3 bidders

	β						
	0.10	0.20	0.50	1	2	10	20
δ	1.0965	1.0889	1.0715	1.053	1.0352	1.0093	1.0048
γ_β	0.9141	0.9202	0.9345	0.9500	0.9664	0.9909	0.9952

Note: The numerical solution was found using fzero in Matlab, assuming $\alpha = 1.1$ and $\gamma_1 = 0.95$. The highlighted values of β correspond to the values used in the experiment.

Table 2: Summary of sessions

Treatment	Sessions	Bidders	Ex. Rate	Subjects	β order	Subject-Auctions
$\beta_{1/0.1/20}^3$	2	3	14	24	1, 0.1, 20	1398
$\beta_{1/0.1/20}^6$	2	6	20	42	1, 0.1, 20	2519
$\beta_{1/20/0.1}^3$	2	3	14	33	1, 20, 0.1	1898
$\beta_{20/1/0.1}^3$	3	3	14	51	20, 1, 0.1	1685

Table 3: Mean difference between bid and value by treatment

$\beta = 1$	$\beta = 0.1$	$\beta = 20$	Bidders	β order
2.79 (6.75)	7.46 (13.67)	1.07 (5.61)	3	1, 0.1, 20
4.07 (10.26)	7.54 (12.96)	0.43 (4.60)	6	1, 0.1, 20
4.04 (10.64)	4.62 (9.94)	2.32 (5.72)	3	1, 20, 0.1
2.37 (7.04)	5.64 (10.61)	1.13 (9.08)	3	20, 1, 0.1

Note: Standard deviations in parentheses.

Table 4: Subjects experiencing losses with β order 1, 0.1, 20

Three Bidders	$\beta = 1$	$\beta = 0.1$	$\beta = 20$
Fraction of Auctions Resulting in Losses	5.4%	9.3%	0.6%
Mean Loss (in ECUs)	9.9	1.5	206.7
Median Loss (in ECUs)	6.0	1.0	140.0
Maximum Loss (in ECUs)	43.0	8.0	400.0
Number of Subjects Who Experience a Loss	11	14	2
Six Bidders	$\beta = 1$	$\beta = 0.1$	$\beta = 20$
Fraction of Auctions Resulting in Losses	4.3%	7.1%	0.4%
Mean Loss (in ECUs)	10.8	1.6	65.0
Median Loss (in ECUs)	6.5	1.1	30.0
Maximum Loss (in ECUs)	88.0	8.8	180.0
Number of Subjects Who Experience a Loss	20	27	4

Table 5: Estimated coefficients from a random effects Tobit model of the effect of β on bids

Treatment	Dependent Variable: Bid-Value											
	$\beta_{1/0.1/20}^3$ (1)	$\beta_{1/0.1/20}^6$ (2)	$\beta_{1/20/0.1}^3$ (3)	$\beta_{20/1/0.1}^3$ (4)	$\beta_{1/0.1/20}^3$ (5)	$\beta_{1/0.1/20}^6$ (6)	$\beta_{1/20/0.1}^3$ (7)	$\beta_{20/1/0.1}^3$ (8)	$\beta_{1/0.1/20}^3$ (9)	$\beta_{1/0.1/20}^6$ (10)	$\beta_{1/20/0.1}^3$ (11)	$\beta_{20/1/0.1}^3$ (12)
$D_{\beta=0.1}$	4.675*** (0.502)	3.469*** (0.422)	0.821* (0.434)	3.173*** (0.484)	7.062*** (1.972)	6.317*** (1.660)	2.419 (2.607)	6.921** (3.314)	5.157** (2.128)	5.410*** (1.767)	-0.638 (2.569)	7.002** (3.231)
$D_{\beta=20}$	-1.613*** (0.508)	-3.634*** (0.417)	-1.638*** (0.430)	-2.427*** (0.474)	-6.098** (3.022)	-5.008** (2.469)	-0.809 (1.697)	-4.441** (1.891)	-7.985** (3.144)	-5.916** (2.538)	-3.972** (1.734)	-3.532** (1.960)
Period					0.043 (0.066)	0.116** (0.055)	0.175*** (0.056)	-0.017 (0.060)	-0.096 (0.086)	0.049 (0.071)	-0.053 (0.069)	-0.018 (0.058)
Period $\times D_{\beta=0.1}$					-0.110 (0.089)	-0.173** (0.075)	-0.172** (0.075)	-0.068 (0.081)	0.030 (0.105)	-0.106 (0.087)	0.056 (0.084)	-0.069 (0.079)
Period $\times D_{\beta=20}$					0.056 (0.088)	-0.066 (0.072)	-0.144* (0.077)	0.180** (0.082)	0.194* (0.104)	0.002 (0.085)	0.087 (0.086)	0.105 (0.094)
Constant	4.246** (1.590)	3.494*** (1.039)	4.593*** (1.103)	3.131* (1.604)	3.800** (1.722)	2.331** (1.176)	2.839** (1.236)	3.728 (2.416)	5.794*** (1.907)	3.169** (1.332)	6.060*** (1.383)	3.528 (2.440)
Observations	1398	2519	1898	1685	1398	2519	1898	1685	1326	2393	1799	1539

* p<0.10, ** p<0.05, *** p<0.01

Note: Standard errors are reported in parentheses. Standard errors are clustered at the individual level. All models include session fixed effects, which have been omitted from the table for brevity. None of session fixed effects are significant. The models in columns 9-12 exclude the first three periods of data in each session. Bids are censored at 100.

Table 6: Comparing estimated log likelihoods across models and auction formats

	SPA	English	FPA
optimistically Irrational	13415	26212	37574
	$\alpha = 1.82, \sigma = 10.18, \gamma = 0.95$	$\alpha = 1.82, \sigma = 10.46, \gamma = 0.95$	$\alpha = 1.78, \sigma = 8.54, \gamma = 0.95$
SNE with normal errors	13687	26480	38036
	$\sigma = 10.99$	$\sigma = 10.87$	$\sigma = 10.87$
SNE with normal errors and JOW	13323	26107	37459
	$\sigma = 9.92, j = 0.77$	$\sigma = 10.30, j = 0.78$	$\sigma = 8.45, j = 0.79$

Note: The SPA models are estimated including data from β order 1, 0.1, 20 for both the 3 and 6 bidder cases. The English auction data come from Georganas (2011) and the group size is 4. The FPA data come from Dyer et al. (1989) with 3 and 6 bidder auctions.

Table 7: Comparing estimated log likelihoods across models, auction formats, values of β , and number of bidders

	SPA		SPA		English		FPA		
	$N = 3$ $\beta = 1$	$N = 3$ $\beta = 0.1$	$N = 3$ $\beta = 20$	$N = 6$ $\beta = 1$	$N = 6$ $\beta = 0.1$	$N = 6$ $\beta = 20$	$N = 4$ $\beta = 1$	$N = 3$ $\beta = 1$	$N = 6$ $\beta = 1$
optimistically Irrational Bidders	1151	2169	1518	2263	3609	2831	13002	5515	5517
SNE with normal errors and JOW	1151	2093	1516	2295	3547	2825	12995	5546	5491

Note: The first three columns correspond to the 3 bidder auctions for the SPA, while the next three columns correspond to the 6 bidder auctions. The English auction data come from Georganas (2011) with 4 bidders per auction. The FPA data come from Dyer et al. (1989). The first column of FPA results come from 3 bidder auctions and the second column come from 6 bidder auctions.

Table 8: Comparing Predicted Mean Overbidding to Observed Mean Overbidding

	SPA			SPA			English		
	$\beta = 1$	$\beta = 0.1$	$\beta = 20$	$\beta = 1$	$\beta = 0.1$	$\beta = 20$	$\beta = 1$	$\beta = 0.1$	$\beta = 20$
optimistically Irrational	2.89	5.12	0.26	3.13	5.40	0.27	0		
SNE with normal errors and JOW	0.78	7.84	0.04	0.78	7.84	0.04	0.78		
Observed	3.26	7.46	1.07	4.48	7.54	0.43	1.17		

Note: The first three columns correspond to the 3 bidder auctions for the SPA, while the next three columns correspond to the 6 bidder auctions. The English auction data come from Georganas (2011) with 4 bidders per auction. The FPA data come from Dyer et al. (1989). The first column of FPA results come from 3 bidder auctions and the second column come from 6 bidder auctions.

References

- [1] Andreoni, James; Yeon-Koo Che & Jinwoo Kim (2007) "Asymmetric Information about Rivals' Types in Standard Auctions: An Experiment," *Games and Economic Behavior*, vol. 59(2), 240-259.
- [2] Breitmoser, Yves (2015) "Knowing Me, Imagining You: Projection and Overbidding in Auctions", *working paper*.
- [3] Camerer, Colin F. & Teck Hua Ho (1999) "Experience-Weighted Attraction Learning in Normal-form Games," *Econometrica*, vol. 67, 827-874.
- [4] Cooper, David J. & Hanming Fang (2008) "Understanding Overbidding in Second Price Auctions: An Experimental Study," *Economic Journal*, vol. 118(532), 1572-1595.
- [5] Crawford, Vincent P. & Nagore Iriberri (2007) "Level-k Auctions: Can a Non-Equilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?," *Econometrica*, vol. 75(6), 1721-1770.
- [6] Dechenaux, Emmanuel; Dan Kovenock & Roman Sheremeta (2015) "A survey of experimental research on contests, all-pay auctions and tournaments," *Experimental Economics*, 18(4), 609-669.
- [7] Dyer, Douglas; John Kagel & Dan Levin (1989) "Resolving Uncertainty About Numbers of Bidders in Independent Private Value Auctions: An Experimental Analysis," *RAND Journal of Economics*, vol. 2, 268-279.

- [8] Fishbacher, Urs (2007) " z-Tree: Zurich Toolbox for Ready-made Economic Experiments, *Experimental Economics*", *Experimental Economics*, vol. 10(2), 171-178.
- [9] Fudenberg, Drew & David Levine, *The Theory of Learning in Games*, 1st edition, MIT Press, Cambridge, MA, 1998.
- [10] Georganas, Sotiris (2011) "English Auctions with Resale: An Experimental Study", *Games and Economic Behavior*, vol. 73(1), 147-166.
- [11] Georganas, Sotiris & Rosemarie Nagel (2011) "English Auctions with Toeholds: An Experimental Study," *International Journal of Industrial Organization*, vol. 29(1), 34-45.
- [12] Goeree, Jacob K.; Charles A. Holt & Thomas R. Palfrey (2002). "Quantal Response Equilibrium and Overbidding in Private-Value Auctions," *Journal of Economic Theory*, vol. 104(1), 247-272.
- [13] Harrison, Glenn (1989) "Theory and Misbehavior of First-Price Auctions", *American Economic Review*, vol. 79(4), 749-762
- [14] Kagel, John H. (1995) Auctions: A Survey of Experimental Results In: Kagel, John H. and Alvin Roth (eds.) *The Handbook of Experimental Economics*, Princeton University Press, Princeton NJ
- [15] Kagel, John H & Dan Levin (1993) "Independent Private Value Auctions: Bidder Behaviour in First-, Second- and Third-Price Auctions with Varying Numbers of Bidders," *Economic Journal*, vol. 103(419), 868-79, July.

- [16] Kagel, John H. ; Ronald M. Harstad & Dan Levin (1987) "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study", *Econometrica*, vol. 55(6), 1275-1304
- [17] Levin, Dan; James Peck & Asen Ivanov (2016) "Separating Bayesian Updating from Non-probabilistic Reasoning: An Experimental Investigation," *American Economic Journal: Microeconomics*, vol. 8(2), 39-60
- [18] Li, Shengwu (2016) "Obviously Strategy-Proof Mechanisms," *SSRN Working Paper 2560028*
- [19] McKelvey, Richard & Thomas R. Palfrey (1995) "Quantal Response Equilibria in Normal Form Games," *Games and Economic Behavior*, vol. 10(1), 6-38.
- [20] Merlob, Brian ; Charles Plott & Yuanjun Zhang (2012) "The CMS Auction: Experimental Studies of a Median-Bid Procurement Auction with Non-Binding Bids", *The Quarterly Journal of Economics*, vol. 127(2), 793-828.
- [21] Noussair, Charles ; Stephane Robin & Bernard Ruffieux (2004) "Revealing Consumers' Willingness-to-Pay: A Comparison of the BDM Mechanism and the Vickrey Auction," *Journal of Economic Psychology*, vol. 25(6), 725-741.
- [22] Ockenfels, Axel & Reinhard Selten (2005) "Impulse balance equilibrium and feedback in first price auctions," *Games and Economic Behavior*, vol. 51 (1), 155-170.

- [23] Roth, Alvin E. & Ido Erev (1998) "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique Mixed-strategy Equilibria," *American Economic Review*, vol. 88, 848-881.
- [24] Selten, Reinhart & Rolf Stoecker (1986) "End Behavior in Sequences of Finite Prisoner's Dilemma Supergames: A Learning Theory Approach," *Journal of Economic Behavior and Organization*, vol. 7, 47-70.
- [25] Vickrey, William (1961) "Counterspeculation and Competitive Sealed Tenders," *Journal of Finance*. Vol. 16(1), 8-37.

A Appendix

A.1 Second Price Auctions

The maximization problem of an optimistically irrational bidder in an SPA in our model is:

$$\max_{b \geq 0} [\sigma(b)]^{\alpha(n-1)} \left\{ x - \left[\frac{[\sigma(x)]^{\alpha(n-1)}}{[\sigma(b)]^{\alpha(n-1)}} \right] \theta_n(x) + \left[\frac{[\sigma(b)]^{\alpha(n-1)} - [\sigma(x)]^{\alpha(n-1)}}{[\sigma(b)]^{\alpha(n-1)}} \right] \gamma(x) \frac{b+x}{2} \right\},$$

where $\theta_n(x)$ denotes the expected price if a bidder wins at a price below his value and $\sigma(b)$ denotes the inverse bidding function.

The first-order condition (FOC) for the maximization problem is:

$$\begin{aligned} 2x\alpha(n-1)[\sigma(b)]^{\alpha(n-1)-1}\sigma'(b) - \gamma(x)\alpha(n-1)[\sigma(b)]^{\alpha(n-1)-1}\sigma'(b)(b+x) \\ - \gamma(x)[[\sigma(b)]^{\alpha(n-1)} - [\sigma(x)]^{\alpha(n-1)}] = 0 \end{aligned}$$

After simplifying we have:

$$\begin{aligned}
2x\alpha(n-1)[\sigma(b)]^{\alpha(n-1)-1}[1-\gamma(x)]\sigma'(b) - \gamma(x)\alpha(n-1)[\sigma(b)]^{\alpha(n-1)-1}(b-x)\sigma'(b) \\
-\gamma(x)[[\sigma(b)]^{\alpha(n-1)} - [\sigma(x)]^{\alpha(n-1)}] = 0
\end{aligned} \tag{A.1}$$

Proof of Proposition 1

Proof. Assume that $\forall x \in [0, 1]$, $b(x) = x$. In such a case equation (A.1) becomes

$$2x\alpha(n-1)[\sigma(b)]^{\alpha(n-1)-1}(1-\gamma(x)) = 0$$

implying $\gamma(x) = 1$, for all $x > 0$, as $\sigma'(b) = 1$ with $b(x) := x$.

Assume that $\forall x \in [0, 1]$, $\gamma(x) := 1$. In such a case equation (A.1) becomes

$$-\{[\alpha(n-1)\sigma(b)]^{\alpha(n-1)-1}(b-x)\sigma'(b) + [[\sigma(b)]^{\alpha(n-1)} - [\sigma(x)]^{\alpha(n-1)}]\} \leq 0 \quad \text{as} \quad b(x) \geq x.$$

□

Derivation of remark 1

Consider the FOC in equation (A.1). Recognizing that at the solution $\sigma(b) = x$ and $\sigma'(b) = \frac{1}{b'(x)}$, and simplifying, we can rewrite (A.1) as

$$2\alpha(n-1)[1-\gamma(x)] - \gamma(x)\alpha(n-1)\left(\frac{b}{x} - 1\right) - \gamma(x)b'(x)\left[1 - \left(\frac{\sigma(x)}{x}\right)^{\alpha(n-1)}\right] = 0. \tag{A.2}$$

Assume that $\forall x \in [0, 1], \gamma(x) = \gamma < 1$ and $b(x) = \delta x$, with $\delta \geq 1$. Equation (A.2) can be written as:

$$2\alpha(n-1)[1-\gamma] - \gamma\alpha(n-1)(\delta-1) - \gamma\delta[1 - (\frac{1}{\delta})^{\alpha(n-1)}] = 0 \quad (\text{A.3})$$

When $\delta = 1$ then (A.3) becomes $2\alpha(n-1)[1-\gamma] > 0$ when $\gamma < 1$, which is the case for overbidding.

By inspecting (A.3), it is clear that the LHS is strictly declining in δ , so that there is a unique δ^* that solves

$$2\alpha(n-1)[1-\gamma] - \gamma\alpha(n-1)(\delta^*-1) - \gamma\delta^*[1 - (\frac{1}{\delta^*})^{\alpha(n-1)}] = 0 \quad (\text{A.4})$$

A.2 First Price Auctions

The maximization problem of an optimistically irrational bidder in an FPA:

$$\max_{B \geq 0} [\sigma(B)]^{\alpha(n-1)} [x - B] \quad (\text{A.5})$$

Denote the solution to (??) by $B(x; \alpha)$ and the risk-neutral Nash equilibrium by $B(x; \hat{\alpha})$. It is well known that the solution to (5) is given by,

$$B(x; \alpha) = x - \frac{\int_0^x t^{\alpha(n-1)} dt}{x^{\alpha(n-1)}} = x - \frac{x}{\alpha(n-1)+1} = \frac{\alpha(n-1)}{\alpha(n-1)+1} x$$

Proposition 2. $B(x; \alpha) \geq B(x; \hat{\alpha})$ as $\alpha \geq \hat{\alpha}$.

Proof. The proof is immediate. □

Thus, optimistically irrational bidders will bid above the RNNE but below their value.

A.3 Joy of Winning

Lemma 3. *In a second price auction where negative payoffs are multiplied times β and bidders have a heterogeneous joy of winning U_i , the dominant strategy equilibrium is $b_i(x_i) = \frac{U_i}{\beta} + x_i$.*

Proof. Expected profits in the auctions are $\Pi_i = \text{prob}\{b_i > \max(b_{-i})\}(x_i + U_i - E[\max(b_{-i})|b_i > \max(b_{-i})])$, and subjects choose a $b_i(x_i)$ to maximize Π_i . Consider bidding $b_0 < b_i(x_i) = \frac{U_i}{\beta} + x_i$, when it matters i.e. you win with $b_i(x_i)$ but lose with b_0 . It also means that the price the winner pays is $p \in [b_0, b_i(x_i)]$. If $p > x_i$, winning earns $U_i - \beta(p - x_i) \geq U_i - \beta((\frac{U_i}{\beta} + x_i) - x_i) = 0$. When $p \leq x_i$, strict positive payoffs are assured. With a similar step we show that bidding $b_i(x_i)$, weakly dominates bidding $b_0 > b_i(x_i)$. \square

A.4 Examining the consistency of joy of winning

The fit of the SNE and the behavioral models is improved when we allow for joy of winning. But is joy of winning a consistent explanation across all different values of β that we have used in the experiments? As we have seen, joy of winning gives a clear prediction for every β . A person who understands the dominant strategy but has joy of winning j , will overbid by exactly j when $\beta = 1$ but will overbid by $10j$ when $\beta = 0.1$. On the other hand she has to overbid by only $j/20$ when $\beta = 20$. In

following Figure 6 we examine if individual players' behavior is consistent with this model.²⁷




Figure 6: Consistency of the joy of winning hypothesis for each individual across the three values of β .

Every dot is a single observation, plotting the average overbidding of a player for two different values of β . In the left panel for example, we have mean overbidding when $\beta = 1$ in the x axis and mean overbidding when $\beta = 0.1$ in the y axis. According to the theory all observations should lie on a straight line through the origin with slope β_x/β_y . What we see is that many observations lie on or close to the line. However there are still some observations, especially in the second and third panel that lie far away from this line. This indicates that although heterogeneous joy of winning could at first sight partly explain overbidding in second price auctions, it is not a fully consistent explanation.

²⁷The bidding results of the 3 bidder and 6 bidder case are pooled, as the number of bidders makes no difference for the overbidding predicted by the joy of winning model.

Table A.1: Model fit in all treatments with β order 1, 0.1, 20. SNE denotes the symmetric Nash equilibrium. -LL is the negative log likelihood, BIC the Bayesian Information Criterion. Note a level k model is the same as the SNE.

3 bidders (1278 observations)						
	SNE (norm)	SNE+joy (norm)	SNE (logit)	SNE (logit+joy)	QRE	QRE+joy
-LL	4861.4	4729.1	4984.7	4944.9	4959.8	4838
BIC	9723.9	9472.5	9976.6	9904.1	9926.8	9690.3
parameters	10.86	$\sigma = 9.66, j = 0.77$	$\lambda = 0.52$	$\lambda = 0.68, j = 2.4$	$\lambda = 0.54$	$\lambda = 0.71, j = 5.8$
6 bidders (2310 observations)						
-LL	8825.1	8593.2	9212	9057.9	9084	8850.9
BIC	17658	17201	18432	18131	18176	17717
parameters	11.06	$\sigma = 10, j = 0.78$	$\lambda = 0.81$	$\lambda = 1.6, j = 2.2$	$\lambda = 0.83$	$\lambda = 1.08, j = 6.3$

A.5 Goodness of fit