On the Effects of Joint Bidding in Independent Private Value Auctions: An Experimental Study

Kirill Chernomaz

San Francisco State University, Department of Economics, 1600 Holloway Ave, HSS 142, San Francisco, CA 94132, Phone: (415) 405-2699, Fax: (415) 338-1057

Abstract

This article explores the effect of a subset of symmetric bidders joining to bid together. Possible applications include mergers, collusion and legal joint-bidding arrangements. The change produces a "strong" party with a more advantageous value distribution than the remaining "weak" bidder(s). The predicted effects include inefficiency, a decrease in the seller's revenue, and higher bidders' payoffs. When the bidders are risk neutral, the members of the strong party are expected to benefit less than the weak bidders. The prediction is reversed when the bidders are sufficiently risk averse. These hypotheses are tested experimentally. The Nash equilibrium with risk aversion is consistent with some features of the data. Contrary to the theory, joint-bidding increases efficiency and the seller's revenue decreases by less than expected. Strong bidders benefit more than weak bidders indicating that the incentives to bid jointly may be greater than hypothesized. Additionally, the experiment assesses the effect of group decision making. A Nash equilibrium prediction for individual-group differences based on differences in risk attitudes is not supported by the data.

JEL Classification: C72, C91, C92, D44

Keywords: Asymmetric Auctions, Independent Private Value, First-price Auctions, Experimental Economics, Group Decision-making, Joint Bidding, Mergers

1. Introduction

Bidder symmetry with respect to possible values for an item is a typical assumption in the auction literature. There are, however, situations when it may be too strong. It is quite common that observable characteristics give clues to auction participants that some of them are likely to value the item more highly than the others. For example, if an industry is characterized by economies of scale, the size of a company can be an observable indicator of its willingness to pay (value) for an item. A bigger company is more likely to have lower costs that translate into higher values. Similarly, a well-established company may have better financing options and, thus, lower costs of capital compared to a new entrant. Maskin and

Email address: kirillch@sfsu.edu (Kirill Chernomaz)

[1]Abbreviations: NE - (Bayesian) Nash equilibrium; RNNE - Nash equilibrium with risk-neutral bidders; RANE - Nash equilibrium with risk-averse bidders; RNBR - best response for risk-neutral bidders; RABR - best response for risk-averse bidders; WILC - Wilcoxon signed-rank test for matched data; MW - Mann-Whitney two-sample test; BIN - Binomial test

Riley (2000) study three classes of 2-bidder first-price auctions where ex-ante asymmetry is common knowledge. They assume that the bidders' value distributions can be ranked so that for one bidder ("strong") high values are more likely than for the other ("weak"). They characterize equilibrium behavior and show that, unlike in second-price auctions where bidding one's value remains the dominant strategy, in first-price auctions asymmetry in value distributions translates into asymmetry of risk-neutral Nash equilibrium (RNNE) bidding functions. Specifically, the weak bidder bids pointwise higher than the strong bidder so that if both have the same value for the item the weak bidder submits a higher bid.

This paper investigates a first-price auction with asymmetries that arise when the value of the strong bidder is the maximum of two independently and identically distributed random draws (signals). One interpretation of this model is that two ex-ante symmetric bidders join to bid together. The members of the new entity no longer compete against each other and submit a single bid. If the joint-bidding entity wins the auction the item is allocated to the member with the higher signal. This model is applicable in various degrees to auctions with collusion, joint bidding, and mergers between independent bidders. As an illustration, consider two firms participating in an auction for a plot of land. One of them has a more profitable opportunity to use the land as it may have more innovative personnel, a better technology, or any other unobservable advantage. If the two firms merge we can expect that the land use proposal offering the highest return will be adopted. Under effective management, the merged firm will adopt the best practices, the best technological innovations, etc. of the two companies. This is exactly equivalent to the maximum operator applied to the firms' values for the auctioned item. Lebrun (1999) proves the existence and uniqueness of the Nash equilibrium for asymmetric auctions of this type. Lebrun (1999) and Maskin and Riley (2000) show that the weak bidder bids pointwise higher than the strong bidder in this setting. Waehrer (1999) derives certain equilibrium properties of this model as it applies to mergers and joint bidding. He points out that in the first-price auction there may be a disincentive to bid jointly since in equilibrium the per-member payoff of the joint-bidding party is expected to be lower than the payoff of the bidders left out of the joint-bidding arrangement.

This paper focuses on the anti-competitive effects of joint bidding. In particular, we compare a 3-bidder symmetric auction to its 2-bidder asymmetric counterpart where two of the three symmetric bidders bid jointly. The main effect of joint bidding is a reduction in competition as it effectively decreases the number of bidders. Consequently, all bidders submit lower bids and auctioned items exchange hands at lower prices bringing lower revenues to the seller and higher profits to the buyers. Furthermore, the resulting asymmetry in equilibrium bidding functions brings about the possibility of inefficient allocations since the weak bidder can win the auction when his or her value is not the highest. We confirm that the weak bidder's payoff increase by more than the per-member payoff of the strong entity but only if the bidders are not too risk-averse. We show that in the Nash equilibrium with moderately risk-averse bidders (RANE) the theoretical result is reversed. In the RANE, the members of the strong party benefit more from the joint-bidding arrangement than the weak bidder especially if we consider changes in the expected utility rather than the expected profit. Is bidders' risk-neutrality an appropriate assumption? Do all bidders decrease their bids when a joint-bidding entity is formed? What is the effect on the seller's revenue and the bidders' profits? Who benefits more: the left-out bidder or the members of the strong party? This paper reports results of an experiment that attempts to answer these questions.

The experimental results support the prediction that in the asymmetric auction strong bidders should bid pointwise lower than weak bidders. In our setup, the asymmetry arises when, as a result of joint bidding, strong bidders lower their bids by a wider margin than do weak bidders. However, in the data we find that strong bidders decrease their bids to a much smaller extent than predicted by either the RNNE or the RANE while the bids of the subjects in the weak role tend to go up. The latter result is inconsistent with the Nash equilibrium regardless of the level of risk aversion. These deviations from the predicted behavior reduce the negative effect of joint bidding on the seller's revenues. Additionally, the level of the seller's revenues in both environments is much higher than in the Nash equilibrium with risk-neutral bidders.

These findings are in line with prior experimental studies that examined the behavior in the asymmetric auctions characterized in Maskin and Riley (2000). Pezanis-Christou (2002) tested a 2-bidder model with uniformly distributed values where the weak bidder has a large positive probability of having the lowest possible value (mass point at the lower bound of the value distribution support). In equilibrium under the first-price auction rule the strong bidder should submit very low bids for a significant interval of values (so-called low-balling). As a result, the first-price auction is expected to bring lower revenues to the seller than the second-price auction. The author finds that subjects in the advantageous position do submit lower bids. However, the bids are not low enough and the revenue in the first-price auction turns out to be higher than in the second-price auction. Guth et al. (2005) report experimental results in a similar model where the value distribution of the strong bidder is a stretched-to-the-right version of his opponent's. In the equilibrium of the first-price auction, the strong bidder should again bid pointwise lower. The strong bidders' optimal degree of bid shading relative to their values is significant enough for their weak opponents to prefer this auction format over the second-price rule. The authors find that in the lab disadvantaged bidders do bid pointwise higher. However, contrary to the theory, weak bidders prefer second-price auctions as the format brings them higher profits. This result is also likely to be driven by the fact that in the first-price auction strong bidders shade their bids less than predicted. Unlike Pezanis-Christou (2002) and Guth et al. (2005) who focus on comparing the performance of different auction rules in the presence of bidder asymmetry, we assess the effect of a change in the auction market structure leading to asymmetry holding the auction rule constant. Furthermore, the asymmetry studied in this paper is more subtle and constitutes a more challenging test of the theory.

We also find that a powerful counteracting force reverses the efficiency implications of joint bidding: the observed efficiency is higher in the asymmetric auction than in the corresponding symmetric environment. One possible explanation for this phenomenon is the heterogeneity in individual bidding behavior and the random variability of bidding patterns. With fewer bidders the distance between the first and second highest values increases so that the probability that a "trembling-hand" bid will lead to an inefficient allocation decreases. By decomposing the observed efficiency we find support for this explanation. Experimental data also suggest that joint bidding does not disproportionately benefit weak bidders as suggested by the RNNE, thus increasing the incentives to bid jointly if the choice were endogenous.[2]

[2] Huck et al. (2007) report findings with a similar flavor in experimental Cournot markets. Contrary to the so-called merger paradox, according to which mergers in Cournot markets are not profitable, they report that

The result is qualitatively consistent with the Nash equilibrium when bidders are moderately risk-averse, which implies that the degree of risk aversion can be an important factor in determining the incentive to merge or bid jointly. Furthermore, we find that the strong bidders' failure to adjust bids downwards to the extent predicted by the RANE makes joint bidding even more attractive to the participating bidders relative to those left out of the arrangement.

A secondary issue investigated in this paper is the effect of the number of decision makers on auction outcomes. There is growing evidence from experimental economics that in some contexts the number of decision makers matters and in particular that individuals and groups behave differently.[3] The few papers on group-individual differences in auction settings suggest that groups may be more competitive. Cox and Hayne (2006) try to disentangle the information pooling effect of joint bidding in common value auctions from the effect of the increasing the number of "heads". While they do not find any difference between (5-person) groups and individuals in the case when information level is low, groups fail to adjust their bids downward as much as individuals when the information level increases. Sutter et al. (2007) examine group and individual behavior in ascending sealed-bid English auctions with common and private value components. They find that groups tend to remain active longer, paying higher prices and receiving lower profits.

Since the type of bidder asymmetry studied here can arise as a result of joint bidding, it provides a natural and convenient framework to study the issue of group-individual differences. We introduce group decision-making into our experiment in a straightforward manner. Initially subjects bid in a symmetric auction, which is followed by the exogenous creation of a joint-bidding entity whose value is the maximum of the values of its members. In this asymmetric market we allow communication between the members of the joint-bidding entity and collective bid generation. We also incorporate a benchmark intermediate stage where each member of the strong entity individually makes a decision on what to bid given the new auction structure. This allows us to separate the effect of changing the auction structure from the effect of changing the number of the decision-makers. One of the factors hypothesized to cause individual-group differences is risk aversion. Groups can be more risk-averse ("risky shift") or less risk-averse ("cautious shift") than individuals. Using the theoretical model we show how a change in the strong bidder's degree of risk aversion is expected to affect equilibrium bidding behavior. In the data, we find no systematic effect of more "heads" on bidding behavior of the strong bidders. In contrast, their weak opponents seem to respond by bidding higher when they know that two people decide on a bid rather than just one. These results cannot be explained by a change in the strong bidder's risk aversion alone.

2. Theoretical Model

We compare two environments: a symmetric auction and the corresponding asymmetric auction following formation of a joint-bidding entity. In the symmetric environment, three bidders compete against each other and bid independently. They are symmetric in the sense that their monetary values for the object can be viewed as drawn from the same distribution.

merged firms are able to obtain higher profits at least in the short run as a result of more aggressive behavior (higher production).

[3] For example, Cooper and Kagel (2005) study groups and individuals in signaling games, Bornstein and Yaniv (1998) in bargaining games, Kocher and Sutter (2005) and Sutter (2005) in beauty contest games.

In the asymmetric environment two of the three bidders do not bid against each other, submit a single bid, and in the case of winning allocate the item to the member with the higher value. Joint bidding effectively creates two bidders with different value distributions. We derive the Bayesian Nash equilibrium (NE) bidding functions for the two cases in turn.

Suppose three bidders participate in a first-price sealed-bid auction. Participants submit sealed bids, which are not revealed until the winner is determined. The winner of the auction is the bidder with the highest bid. The winner pays her bid and receives the item. Each bidder ranks monetary outcomes according to the same constant relative risk aversion (CRRA) utility function $u(x) = \frac{x^{1-r}}{1-r}$, $0 \leq r < 1$. Their monetary values for the object are drawn independently from the uniform distribution on the interval $[0, \bar{v}]$ with a cumulative density function (cdf) $F(v) = \frac{1}{\bar{v}}v$. Assuming that her opponents use the same monotone-increasing bidding function $\beta(v)$ with an inverse $\sigma(b)$, by submitting a bid b_1 bidder 1 wins the auction with probability $P(b_1 > \max\{b_2, b_3\}) = P(\sigma(b_1) > \max\{v_2, v_3\}) = [F(\sigma(b_1))]^2$. The other two bidders are in exactly the same situation. Thus, the objective of bidder i is to choose a bid b to solve the following problem:

$$\max_b u(v_i - b)[F(\sigma(b))]^2$$

The first-order condition is:

$$-u'(v_i - b)[F(\sigma(b))]^2 + 2u(v_i - b) F(\sigma(b)) F'(\sigma(b)) \sigma'(b) = 0$$

Simplifying and noting that in a symmetric equilibrium it must be that $v_i = \sigma(b)$ we get:

$$2u(\sigma(b) - b) F'(\sigma(b)) \sigma'(b) = u'(\sigma(b) - b) F(\sigma(b))$$

Substituting in the functional forms of the utility and the probability distribution functions we obtain:

$$\frac{2}{1-r}\sigma'(b) b + \sigma(b) = \frac{2}{1-r}\sigma'(b)\sigma(b)$$

Multiplying both sides by $\sigma(b)^{\frac{2}{(1-r)}-1}$ the first-order condition can be written as:

$$\frac{d}{db}\left\{[\sigma(b)]^{\frac{2}{1-r}} b\right\} = \frac{2}{1-r}\sigma'(b)\sigma(b)^{\frac{2}{1-r}}$$

Integrating and noting that the constant of integration is 0 (due to the initial condition $\sigma(0) = 0$) we obtain:

$$\begin{aligned}[\sigma(b)]^{\frac{2}{1-r}} b &= \frac{2}{1-r}\int_0^b \sigma'(s)\sigma(s)^{\frac{2}{1-r}} ds \\ &= \frac{2}{3-r}\sigma(b)^{\frac{3-r}{1-r}}\end{aligned}$$

The candidate inverse equilibrium bidding function is then:

$$\sigma(b) = \frac{3-r}{2} b$$

and the corresponding candidate equilibrium bidding function is:

$$\beta(v) = \frac{2}{3-r}v$$

Note that $\beta(v)$ is monotone-increasing validating our original assumption. It can be shown that monotonicity also ensures that the second order condition is satisfied.[4] Thus, $\beta(v)$ is indeed a symmetric Bayesian Nash Equilibrium bidding function. It tells us that risk-neutral bidders ($r = 0$) would use $\beta(v) = \frac{2}{3}v$ in the equilibrium (RNNE). Risk-averse bidders ($0 < r < 1$) would bid pointwise higher (RANE). Analogous calculations yield a solution for an arbitrary number of bidders N: $\beta_N(v) = \frac{N-1}{N-r}v$.[5]

Now, suppose that two out of the three bidders decide to form a joint-bidding entity: they submit a single bid and in the case of winning allocate the object to the member with the higher value. Thus, the entity's value is the maximum of the values of its members. Since the cdf of the maximum of independent random variables on the same interval is the product of their cdf's, the value distribution of the joint-bidding entity is $F_s(v) \equiv (F(v))^2 = \frac{1}{\bar{v}^2}v^2$. The value distribution of the remaining bidder is the same as in the symmetric environment which we now denote as $F_w(v) \equiv F(v) = \frac{1}{\bar{v}}v$. In essence, joint bidding leads to a reduction in the number of bidders and creates a party whose value distribution is more favorable compared to the other bidder in the sense of stochastic dominance. We shall often refer to the joint-bidding entity as the strong bidder and to the other participant as the weak bidder (hence the subscript notation).

Assuming that bidder i's opponent uses a monotone-increasing bidding function $\beta_j(v)$ ($i, j \in \{s, w\}$, $i \neq j$) with an inverse $\sigma_j(b)$, by submitting a bid b bidder i wins the auction with probability $P(b_j < b) = P(v_j < \sigma_j(b)) = F_j(\sigma_j(b))$. We shall further assume that the two bidders may differ in their risk attitude, i.e. their utility function is $u_i(x) = \frac{x^{1-r_i}}{1-r_i}$. Thus, bidder i faces the following problem:

$$\max_b u_i(v_i - b) F_j(\sigma_j(b)), \; i, j \in \{s, w\}, \; i \neq j$$

The first order condition is:

$$-u_i'(v_i - b) F_j(\sigma_j(b)) + u_i(v_i - b) F_j'(\sigma_j(b)) \sigma_j'(b) = 0$$

which after imposing the equilibrium condition $v_i = \sigma_i(b)$, substituting in the utility function, and rearranging becomes:

$$\begin{aligned} \sigma_j'(b) &= \frac{(1-r_i)}{(\sigma_i(b) - b)} \frac{F_j(\sigma_j(b))}{F_j'(\sigma_j(b))} \\ &= \frac{1}{(\sigma_i(b) - b)} \frac{\tilde{F}_j(\sigma_j(b))}{\tilde{F}_j'(\sigma_j(b))} \end{aligned} \tag{1}$$

where $\tilde{F}_j(v) \equiv [F_j(v)]^{\frac{1}{(1-r_i)}}$. It can be quickly verified that the differential equation in (1) is the same as one would obtain if the bidders were risk-neutral with values distributed according

[4] For example, see Holt (1980).

[5] See Cox et al. (1982) and Cox et al. (1988) for a partial solution of the problem allowing for heterogeneous risk aversion.

to the modified probability distribution functions $\tilde{F}_j(v)$. For the latter class of problems Lebrun (1999) proves that a unique Bayesian Nash Equilibrium exists as long as $\frac{\tilde{F}_s(v)}{\tilde{F}_w(v)}$ is strictly increasing on the interval $(0, \bar{v}]$ (the conditional stochastic dominance assumption $\frac{\tilde{F}_s'(v)}{\tilde{F}_s(v)} > \frac{\tilde{F}_w'(v)}{\tilde{F}_w(v)}$ in Maskin and Riley (2000)). In which case, the equilibrium bidding strategies $\beta_s(v)$ and $\beta_w(v)$ are monotone increasing so that their inverses $\sigma_s(b)$ and $\sigma_w(b)$ exist on the interval $[0, \eta]$, where $\eta \in (0, \bar{v})$. In our model, this equilibrium requirement amounts to:

$$\frac{d}{dv}\left[\frac{\tilde{F}_s(v)}{\tilde{F}_w(v)}\right] = \frac{d}{dv}\left[\left(\frac{1}{\bar{v}}v\right)^{\frac{2}{(1-r_w)}-\frac{1}{(1-r_s)}}\right] > 0, \forall v \in (0, \bar{v}]$$
$$\Leftrightarrow 2(1-r_s) > (1-r_w). \quad (2)$$

Thus, as long as the joint-bidding entity is not much more risk-averse than the remaining bidder the FOC in (1) characterize the unique equilibrium.

Substituting the value distribution functions into (1) we obtain:

$$\begin{cases} \sigma_s'(b) = \frac{(1-r_w)}{(\sigma_w(b)-b)}\frac{\sigma_s(b)}{2} \\ \sigma_w'(b) = \frac{(1-r_s)}{(\sigma_s(b)-b)}\sigma_w(b) \end{cases} \quad (3)$$

The equilibrium strategies $\beta_s(v)$ and $\beta_w(v)$ can be obtained as inverses of the solution to the system of differential equations formed by the first-order conditions (3), and the two initial conditions:

$$\sigma_i(0) = 0,\ i \in \{s, w\}$$
$$\sigma_i(\eta) = \bar{v}.$$

The first initial condition states that both bidders must bid 0 whenever they receive the lowest possible value of 0 (by individual rationality and implicitly assuming that the seller's reservation price is 0). The second condition requires that the bid associated with the highest value should be the same for both bidders. Given that one bidder's highest possible bid is η, rationality requires that the other bidder does not bid above η. A symmetric argument establishes that the highest possible bids must be the same for both bidders.

Lebrun (1999) and Maskin and Riley (2000) prove that the equilibrium bidding functions satisfy the following condition: $\beta_s(v) < \beta_w(v)$ for all $v \in (0, \bar{v})$ as long as the stochastic dominance condition is satisfied, which in our setting requires that (2) hold. In other words, a weak bidder with value v always submits a higher bid than a strong bidder with the same value in the equilibrium, as long as the strong bidder is not too risk-averse relative to the weak bidder. We refer to this result as the weak bidder bidding *pointwise higher* than the strong bidder. Additionally, under the same condition the resulting bid distribution of the strong bidder first-order stochastically dominates the bid distribution of the weak bidder: $F_{\beta_s}(b) < F_{\beta_w}(b)\ \forall b \in (0, \eta)$, where $F_{\beta_s}(b)$ and $F_{\beta_w}(b)$ are the corresponding bid distributions.

The exact equilibrium bidding functions are obtained numerically.[6] The left panel of Figure 1 plots the equilibrium bidding functions for risk-neutral bidders in both the symmetric

[6] The system of differential equations is solved using the so-called backward-shooting method (see Marshall et al. (1994)) which involves starting from the boundary condition with an initial guess of the common

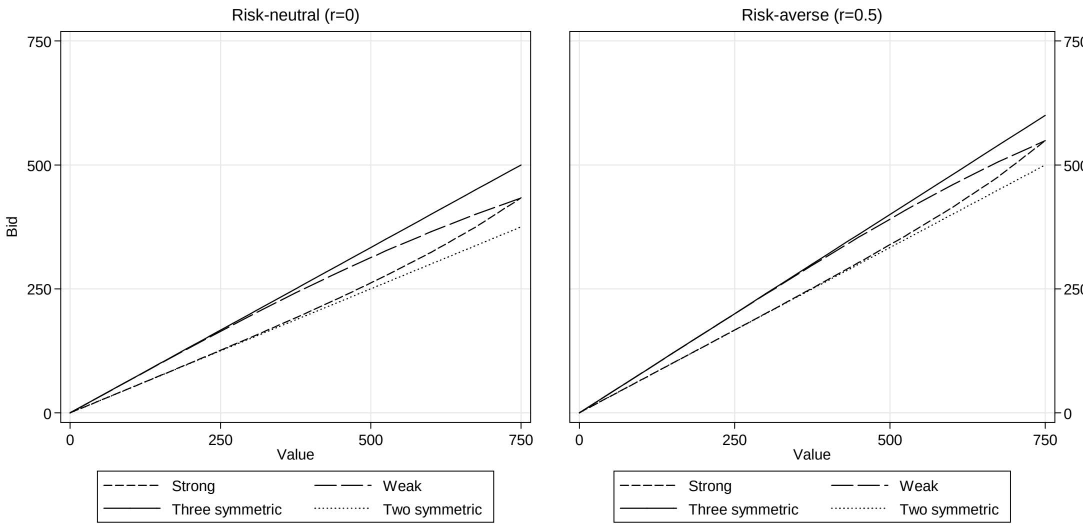

Figure 1: Equilibrium bidding functions: risk-neutral and risk-averse bidders.

and the asymmetric environment using 750 as the upper bound for the value distributions. Another benchmark is also provided in the panel: the equilibrium for the symmetric market with only two bidders. This benchmark helps us to assess the effect of decreasing the number of bidders without changing their value distributions. The figure shows that bidding functions in the asymmetric market are enveloped by the bidding functions in the symmetric markets. Holding the value distribution constant, a decrease in competition encourages the participants to submit lower bids. However, due to his disadvantaged position, the weak bidder has an incentive to bid higher than if his opponent had the same value distribution. In equilibrium both bidders find it optimal to bids pointwise higher than in the symmetric 2-bidder market but pointwise lower than in the symmetric 3-bidder market. The right panel of Figure 1 plots the equilibrium strategies when the bidders are risk-averse with $r_i = 0.5$. It can be seen that risk aversion induces higher bidding in both the symmetric and the asymmetric settings.

Another interesting comparative statics result is illustrated in Figure 2. Each panel depicts two pairs of functions. One pair (solid lines) corresponds to the equilibrium bidding functions of the bidders when they are risk-averse with $r_s = r_w = 0.2$. The other pair (dashed lines) are the equilibrium bidding functions when the risk aversion parameter of the strong bidder changes. In the figure, no distinction is made between the functions of the strong and the weak bidders since the former is always the lower one in each pair of functions. The left panel of Figure 2 shows that a decrease in the degree of risk aversion of the strong bidder leads to a pointwise reduction in equilibrium bids of *both* the strong and the weak bidders.

---

highest bid $\tilde{\eta}$, solving the system backwards using a standard algorithm (we use ode45 Matlab solver), and then verifying whether the initial condition is satisfied. If not, $\tilde{\eta}$ is adjusted in the appropriate direction until the initial condition is satisfied within a specified margin of error. Backward shooting is necessary due to a singularity at the initial condition $\sigma_i(0) = 0$. This method is employed in the BIDCOMP^2 program by Li and Riley (2007) that we used to verify some of our numerical calculations.

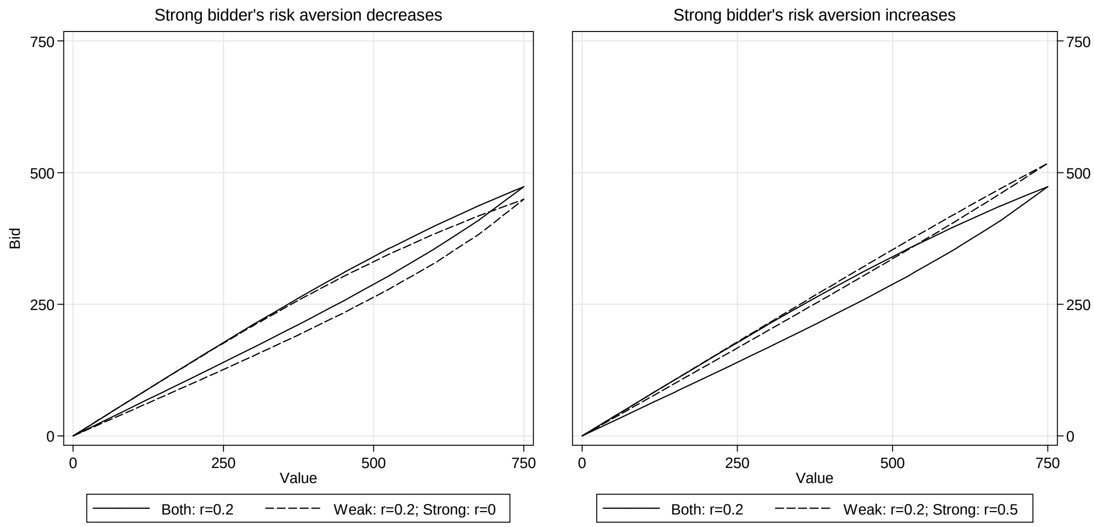

Figure 2: The effect of a change in the strong bidder's risk aversion.

On the other hand, an increase in the strong bidder's risk aversion parameter makes *both* bidders bid higher in equilibrium (see the right panel). It also reduces asymmetry between the bidders. When $r_w = 0.2$ and $r_s = 0.5$ the equilibrium bidding functions come close to being symmetric. In fact, when (2) holds with equality, the degree of risk aversion of the strong bidder exactly offsets his value distribution advantage so that the resulting equilibrium functions are symmetric. The effect of changes in the risk attitude of the strong bidder provides theoretical foundation for the effect of group decision-making. If group decision-making alters the risk attitude of the strong bidder, then we should expect all bidders to change their bids in the same direction. We expect higher bids if the strong bidder becomes more risk-averse, and lower bids if the strong bidder becomes less risk-averse.

Table 1 shows changes in key market variables under the assumption that the upper bound for value distribution is 750.[7] First of all, joint bidding produces inefficiency. Due to the asymmetry of the equilibrium bidding functions inefficient allocations occur with positive probability. An item can go to the bidder who does not have the highest value. Since the weak bidder bids higher for any given value, it is possible that when the values of the two bidders are close the weak bidder can win the item when his value is lower. Thus, the total surplus in the market decreases in the presence of joint bidding. When the bidders are risk-neutral, the predicted inefficiency is around 9 percent. When the bidders are risk-averse with $r_s = r_w = 0.5$, the expected inefficiency is 7 percent as an increase in the level of risk aversion

[7] These results are obtained by simulating 1 million auctions where bidders follow equilibrium bidding strategies and values are independently drawn from the corresponding probability distributions. Since the draws are independent and the variances of all variables are finite, sample averages converge to expected values *almost surely*. The (inverse) bidding strategies are computed numerically for 1000 equally spaced nodes. Continuous approximations are subsequently obtained by fitting 7^{th} order polynomials using OLS ($R^2 = 1.0000000$).

Table 1: Expected equilibrium outcomes

	Efficiency	Revenue	Bidder's profit (utility)	
			Strong per member	Weak
Sym., $r = 0$	100%	375.0	62.5	62.5
Asym., $r_s = r_w = 0$	91.1%	319.7	75.4	88.7
% Change	↓8.9%	↓14.8%	↑20.6%	↑41.9%
Sym., $r = 0.5$	100%	450.0	37.5 (7.0)	37.5 (7.0)
Asym., $r_s = r_w = 0.5$	93.0%	405.6	51.4 (11.0)	52.0 (9.0)
% Change	↓7.0%	↓9.9%	↑37.0% (↑57.2%)	↑38.6% (↑28.5%)

reduces asymmetry in bidding.

Second, the expected price at which the item is sold decreases. Since joint bidding results in a smaller number of bidders, the competition among them lessens. As a result, there is a redistribution of the remaining surplus from the seller to the bidders. The decline in the seller's revenue is not as dramatic as would be the case if one of the bidders simply left the market. As illustrated in Figure 1, both categories of bidders decrease their bids with asymmetry, but the decrease falls short of the level that would be observed in a 2-bidder symmetric market. When the bidders are risk-neutral, the expected decline in the seller's revenue amounts to almost 15 percent: from 375 in a symmetric 3-bidder market to around 320 with joint bidding, versus 250 in a symmetric 2-bidder situation (a drop of 1/3). Risk aversion compels the bidders to submit higher bids in both the symmetric and the asymmetric environments increasing the revenue. Additionally, the decrease in the revenue associated with joint bidding in this case is smaller at around 10 percent.

The flip side of the decrease in the seller's revenue is that all bidders benefit from the decrease in competition. The profits of the strong bidders increase even if we account for sharing of profits. Assuming that members of the strong party split the profits equally, each of them enjoys higher profits compared to the symmetric 3-bidder situation. As expected, with risk neutrality the profits of the weak bidder increase by a substantially higher percentage. While the per-member profits of the strong entity increase by about 21 percent, the weak bidder's profits increase by about 42 percent. Waehrer (1999) suggests that this may create a free-rider problem if the choice to bid jointly is endogenized. A bidder would have an ex-ante disincentive to be a part of a joint-bidding arrangement. He would prefer the others to team up, which would let him enjoy a higher increase in expected payoff. We find, however, that when bidders are moderately risk-averse ($r_s = r_w = 0.5$) the disparity in earnings between the members of the strong entity and the weak bidder essentially disappears. Furthermore, for higher levels of risk aversion the prediction that weak bidders benefit more is reversed (not reported). If, more appropriately, we look at changes in the expected utility (in brackets), the prediction is reversed even at $r_s = r_w = 0.5$. The dramatic increase in the strong bidders' expected utility is due to a reduction in risk since their probability of winning almost doubles (not reported). Thus, when bidders are moderately risk-averse the incentives to be part of a joint-bidding arrangement are stronger as the members benefit more than the left-out bidder.

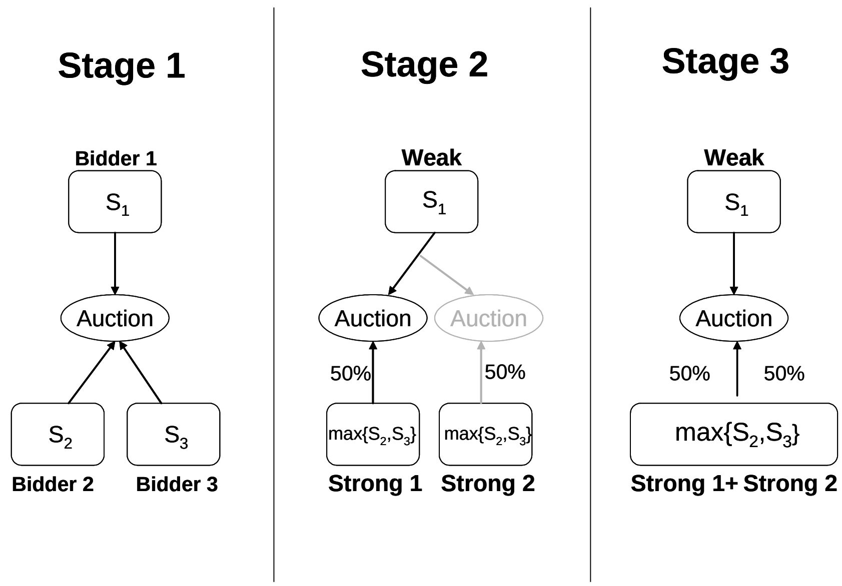

Figure 3: 3-stage experimental design

3. Experimental Design

We use experimental approach to evaluate the hypotheses suggested by the underlying theory. The experiment is designed to compare three environments: 3-bidder symmetric markets, 2-bidder asymmetric markets with a single decision-maker responsible for the strong entity's bid, and 2-bidder asymmetric markets with two decision-makers generating the strong entity's bid. Comparison of the first two environments puts to the test the Nash equilibrium predictions with respect to changes in bidding and market outcomes following creation of a joint-bidding team. Comparison of the last two environments allows us to evaluate the behavioral hypothesis that there are differences in decision-making between individuals and groups. We employ a 3-stage within-subject experimental setup (Figure 3). In this setup subjects make three sequential decisions. First, subjects submit bids in the symmetric environment (stage 1). Then, a 2-bidder joint-bidding arrangement is simulated by combining values of a pair of bidders so that the higher value in the pair becomes the value for both of them, while the value of the third bidder remains unchanged. The two members of the strong entity are first asked to submit new bids without consulting each other (stage 2) and subsequently are allowed to communicate and submit a common bid (stage 3). The third bidder is allowed to submit new bids in stage 2 and stage 3. The advantage of this setup is that it eliminates some of the random noise associated with between-subject designs as the decisions are made by the same individuals based on the same set of values.

The implementation of this 3-stage design was as follows. In the beginning of a session subjects were randomly assigned roles. A third of the subjects were assigned the role of weak bidders and the rest were assigned the role of strong bidders.[8] These roles were assigned

[8] The terms "strong" and "weak" were never used during the experiments. Instead the subjects in the role of weak bidders were referred to as "individuals", and subjects in the role of strong bidders – as members of a "team".

randomly and did not change during the session. The strong bidders were further randomly paired into teams. These two-subject teams did not change either. As a result, a given session had an equal number of teams in the role of strong bidders and individual subjects in the role of weak bidders. Subjects went through 24 periods of auction bidding. In each period teams and individual subjects were randomly matched into 3-person groups so that each group consisted of two strong subjects from the same team and a single weak subject. These groups constituted independent auction markets where bidding took place.

Prior to bidding subjects received *iid* signals S_i from a uniform distribution on the interval [0, 750]. In stage 1 each subject in a group submitted a sealed bid based on her private signal, which was explained to be the value of the object to the bidder.[9] This setup corresponds to the symmetric 3-bidder model. Subjects' roles had no bearing in this stage.[10] Once a bidder had submitted her stage 1 bid she proceeded to stage 2. In this stage strong bidders in the same team were informed about each other's signals. The higher of the two signals was explained to be the new value of the object to both of them (team value). No additional information such as opponent bids, signals, etc., was revealed at that point. Each team member independently submitted a bid against the weak bidder, and the bids were not revealed to their partners. Thus, stage 2 serves as a benchmark where after a change in auction structure bidding decisions are made by individual strong decision-makers. Weak bidders also submitted a stage 2 bid. Weak bidders' values did not change but they were informed that they faced only one other bidder whose value was the highest of the two team members' signals. The profit of each team member was evaluated based on her own bid and the bid of the weak bidder. Therefore, for a team member the partner's bid in stage 2 had no payoff implications. A team member's profit in this stage, as well as in stage 3, was half of the difference between the team value and her bid. This implemented a fifty-fifty profit split by the joint-bidding entity. The profit of the weak bidder was evaluated based on her bid and one of the two bids (chosen randomly) of the strong team members in her group. In stage 3 team members were allowed to communicate. It was implemented using an instant messaging system. Neither category of bidders received any new information in stage 3. The only difference was the open communication link between the strong subjects in the same team. In this stage team members were encouraged to discuss their bidding strategy and submit a common bid. Common bids were not required. Team members could submit different bids in stage 3, in which case one of them was chosen randomly to become the team bid. Weak bidders were asked to submit a stage 3 bid as well. The profits of all three subjects in a group were calculated based on the common team bid and the bid of the weak subject. Team profit was split equally between the team members. At the end of the period full information feedback was provided including the bids and values of all the bidders in one's group.

Three experimental sessions were conducted (see Table 11 in the Appendix). Each session lasted for 2 hours on average. The experiment was computer-based with the software designed using the zTree experimental toolkit (see Fischbacher (2007)). Additionally, proprietary instant messenger software was used to implement the communication. Subjects were

[9] Auctions were framed in terms of resale auctions where signals/values were called "resale prices". The subjects were told that the high bidder pays her bid, obtains the item and automatically resells it to the experimenter at the resale price.

[10] Team members also knew that their partners were bidding against them in the same group.

recruited via bulk e-mail from undergraduate students at the Ohio State University who were enrolled in an economics course.[11] For each session a number of subjects were gathered in the computer lab. Written instructions were distributed and read aloud by the experimenter. Two practice rounds familiarized subjects with the software followed by 24 periods of auction bidding for cash.

Monetary payoffs were cumulative and based on a randomly chosen stage in each period. The stage chosen applied to all subjects. This payoff method was chosen to avoid interdependence of bidding decisions across stages resulting from possible hedging-like considerations on the part of the subjects. The profits from a randomly chosen stage were converted into US dollars at the rate of $1 per 40 experimental currency units (ECU). The accumulated profits were paid to subjects in cash at the end of a session. Earnings averaged $27 per subject including a fixed show-up fee of $6.

3.1. Methodological issues

We chose to employ a within-subject design in order to reduce the noise in the data and increase the power of statistical tests. In between-subject experiments different subjects are used in each treatment. By contrast, in within-subjects experiments the same subject is exposed to two or more treatment conditions. The advantage of the latter approach is that the noise in the data due to unobserved heterogeneity in subjects' individual characteristics is reduced. Any difference between treatments is no longer attributable to random differences in subject samples. As a result, a false null hypothesis that there is no difference between treatment conditions is more readily rejected as more powerful tests can be used. In private information games such as auctions, an additional source of noise is due to randomly drawn bidder values (or, more generally, types). Thus, any observed difference between treatments could also be due to a difference in the samples of values. Although it is possible to use the same sample of values in between-subject experiments it is difficult to implement. In contrast, implementation is straightforward using a within-subject design and allows further reducing the noise in the data. Additionally, asking subjects to submit a sequence of bids in similar conditions encourages them to think about the direction of the change rather than reassessing the absolute magnitude of bids, thus reducing the room for random error. We conducted a couple of pilot experiments using a between-subject version of the design and found that substantial heterogeneity of subjects' bidding behavior is likely to mask the treatment effect under investigation.

The downside of within-subject designs is the order (or carryover) effect. When a subject is exposed to multiple treatments their ordering may produce an effect on its own. The order effect complicates assessment of the treatment effect. A subject's performance in treatment B can be different depending on whether the sequence of treatments is AB or BA. In economics experiments that evaluate subjects' behavior in complex environments the most important factor responsible for the order effect is arguably learning. A subject may learn something about the task after completing the first treatment which may "improve" her performance in the subsequent treatment. Lack of feedback between treatments does not necessarily solve the problem. Weber (2003) reports that in a competitive guessing game convergence to

[11] The pool of potential participants was very diverse since most students in the pool were taking an introductory economics course to fulfill a general education requirement.

the equilibrium is observed even if no feedback is provided between periods. The effect of learning is of particular concern when e.g. the first half of an experimental session is carried out under one treatment condition while in the second half another treatment is used. In our experiment the three treatment conditions are repeated multiple times in the same order with no feedback between treatments and full feedback after each three-treatment period. Multiple repetitions work to reduce the importance of order effects due to learning. Learning between treatments within a period is expected during the trial periods when practice with the software and improved familiarity with the setup can still contribute to differences between adjacent decisions. However, such within-period learning effect is not expected to persist. Within-period learning should be contrasted with between-period learning which is expected to affect all treatments in a similar way. Similar repeated within-subject designs have been employed in, for example, Andreoni et al. (2007) and Kagel and Levin (1993).

A possible remedy for the order effect is counterbalancing: varying the order of treatments. Similar to Andreoni et al. (2007), in this paper the order is determined by the revelation of information and cannot be reversed without modifying some key features of the setup. In stage 2 team members learn each other's signals and in stage 3 they learn each other's views on bidding. This information cannot be "unlearned". Thus, varying the order of treatments is problematic. Alternatively, several periods from the beginning of the experiment can be omitted under the assumption that between-treatment learning continues in the beginning of the experiment but loses its importance towards the end. Along these lines we report changes in bidding behavior over time. However, we generally present our results based on all the available data due to a relatively small number of bidding periods and lack of consistent time patterns across bidder roles, stages, and experimental sessions.

4. Data Analysis

4.1. Bidding

We begin our analysis by examining the observed bidding behavior. Theory predicts that in stage 1 all bidders should bid symmetrically regardless of their role since their role has no bearing on the distribution of values. Nash equilibrium bidding functions in stage 1 are linear with a slope of 2/3 for risk-neutral bidders and steeper slopes if bidders are risk-averse (see Figure 1). Both stage 2 and 3 correspond to the asymmetric auction model with two bidders. The key prediction for these stages is that weak bidders should bid higher than strong bidders at every value. With respect to differences between stages, we expect both bidder categories to decrease their bids between stage 1 and stage 2. If groups behave differently than individuals and if the difference is due to changes in risk attitudes we expect bids to change in the same direction for both bidder categories between stages 2 and 3.

4.1.1. The level of bidding

First, we analyze whether the point predictions of the theory are supported by the data. We find that, consistent with prior research, the RNNE has very weak quantitative predictive power in our data: significant overbidding relative to the RNNE is observed for both bidder categories in every stage. Table 2 reports average bids by intervals of values. The average is taken over all possible values as well as over 8 intervals: 0-50 (interval 0), 50-150 (interval 1), 150-250 (interval 2), ..., 650-750 (interval 7). Average bids implied by the Nash equilibrium bidding with risk-neutral and risk-averse bidders are provided in Table 3. A comparison of

Table 2: Average bids by value intervals.

Value	Stage 1			Stage 2			Stage 3		
Intervals	Strong	Weak	Diff.	Strong	Weak	Diff.	Strong	Weak	Diff.
0: 0-50	16.1	18.2	-2.1	38.0	16.9	21.1	2.0	17.6	-15.6‡
1: 50-150	74.5	79.3	-4.8	87.1	79.9	7.2	82.3	80.3	1.9
2: 150-250	158.3	157.2	1.1	151.2	158.6	-7.4	165.2	158.8	6.3
3: 250-350	231.5	253.5	-22.0	222.6	253.8	-31.2‡	226.0	256.0	-30.0†
4: 350-450	318.1	334.9	-16.8	308.6	332.0	-23.5	314.1	337.0	-22.9
5: 450-550	391.7	412.0	-20.3	386.6	419.2	-32.6‡	380.9	425.2	-44.3♯
6: 550-650	470.1	490.4	-20.3	459.6	492.4	-32.7‡	452.8	493.3	-40.5‡
7: 650-750	524.2	543.4	-19.2	503.1	551.3	-48.3†	502.2	567.5	-65.4♯
Overall	286.8	294.9	-8.1	371.0	297.1	73.9♯	370.4	301.3	69.1♯

† - p<0.1, ‡ - p<0.05, ♯ - p<0.01, Mann-Whitney equality of means test.

the two tables suggests that observed bidding is consistently higher than that implied by the RNNE. This observation is confirmed by the Wilcoxon signed-rank test for matched data (WILC).[12] The null hypothesis that the observed average bids do not differ from the RNNE average bids is rejected in favor of a two-tail alternative (not reported in the table, p<0.01 for all intervals except for the value interval 0). Bidding above the RNNE is a common finding in auction experiments (see Kagel (1995) for a review). Generally, risk aversion can account for such behavior as shown in Figure 1.[13]

Other studies report estimates of the CRRA risk aversion parameter in auction contexts. Those estimates vary significantly. Chen and Plott (1998) obtain estimates between 0.287 and 0.65 using a model that allows for individual heterogeneity in risk aversion. Goeree et al. (2002) combine CRRA utility with the quantal response model of noisy bidding and obtain estimates between 0.51 and 0.56. Estimating the parameter precisely is not the focus of this paper. Instead, we show that while $r = 0.5$ cannot be rejected in the symmetric stage it is not a good descriptor of the behavior in the asymmetric stages. We denote the Nash equilibrium when bidders are CRRA risk-averse with $r = 0.5$ as the $\text{RANE}_{0.5}$. In stage 1 the null hypothesis that average bids do not differ from the $\text{RANE}_{0.5}$ can only be rejected for team members and only for the highest interval of values, 650-750 (p<0.1, WILC, 2-tail). However, in stages 2 and 3 strong bidders bid mostly higher than the $\text{RANE}_{0.5}$ (p<0.1, WILC, 2-tail, value intervals 1-6). Weak bidders bid mostly higher in stage 3 (p<0.1, WILC, 2-tail, value intervals 1, 3-5, 7), but only for interval 5 in stage 2 (p<0.1, WILC, 2-tail). These results

[12] In this test as well as all the other tests we use subject averages or other subject summary measures as a unit of observation. Specifically, instead of 24 observations for a subject we use only one, e.g. average subject's bid. This is done to avoid repeated measure issues.

[13] Typically, models allowing for risk aversion do not capture all aspects of the data (Cox et al. (1988)). A number of alternative explanations have been proposed. Cox et al. (1988) use utility of the event of winning and threshold utility of surplus to account for overbidding and other features of the observed behavior. Goeree et al. (2002) employ the notion of quantal response equilibrium to show that noisy best response in addition to risk aversion offers a superior fit to experimental data. Overbidding can also be explained using a level-k non-equilibrium model (Crawford and Iriberri (2007)) and regret from losing at an affordable price (Filiz-Ozbay and Ozbay (2007)).

Table 3: Average Nash equilibrium bids by value intervals.

Value Intervals	Symmetric			Asymmetric					
	r=0	r=0.5	r=0.8	$r_s=r_w=0$			$r_s=r_w=0.5$		
				Strong	Weak	Diff	Strong	Weak	Diff
0: 0-50	15.1	18.1	20.6	24.4	15.8	8.6	32.6	19.0	13.6
1: 50-150	65.6	78.7	89.4	56.0	64.7	-8.7	74.5	77.7	-3.2
2: 150-250	132.7	159.2	180.9	105.8	130.2	-24.3	140.4	157.3	-16.9
3: 250-350	201.1	241.3	274.2	150.2	197.0	-46.8	198.2	240.4	-42.2
4: 350-450	264.4	317.2	360.5	207.6	255.6	-48.0	271.4	315.4	-44.0
5: 450-550	335.8	403.0	457.9	264.1	317.0	-52.9	341.5	396.0	-54.5
6: 550-650	401.0	481.2	546.8	324.7	366.6	-41.9	415.4	461.7	-46.4
7: 650-750	466.5	559.8	636.2	394.4	411.9	-17.4	500.4	521.1	-20.7
Overall	245.3	294.3	334.4	266.9	226.6	40.3	343.1	282.0	61.1

suggest that the $RANE_{0.5}$ does not fully explain the observed comparative statics results. The level of risk aversion that explains the behavior in the symmetric stage does not do as well in the asymmetric stages. In particular, strong bidders bid higher than predicted by the $RANE_{0.5}$ in the asymmetric stages. To put it differently, they do not reduce their bids sufficiently following the change to the asymmetric auction structure.

Observation 1. *Both bidder categories bid higher relative to the RNNE in all stages. CRRA risk aversion with r=0.5 is a good descriptor of bidding behavior in the symmetric stage. However, it does not explain well the level of bidding in the asymmetric stages.*

4.1.2. Differences between strong and weak bidders

The next hypotheses we test are related to the asymmetry in bidding functions in stages 2 and 3. First, in these stages weak bidders are predicted to bid pointwise higher than strong bidders, i.e. $\beta_s(v) < \beta_w(v)\ \forall v \in (0, \bar{v})$. Second, despite the higher weak bids, the strong bidder's bid distribution is predicted to first-order stochastically dominate that of the weak bidder. The dominance result also implies that the mean weak bid is *lower* than the mean strong bid, $E[\beta_s] > E[\beta_w]$ (Hadar and Russell (1969)). The second prediction can be evaluated by comparing the overall sample averages of strong bids ($\bar{b}_s$) and weak bids ($\bar{b}_w$). The last row of Table 2 reports the overall bid averages for every stage distinguishing between bidders in the two roles. Additionally, the difference between average bids of the two bidder categories is reported ("$Diff.$" $\equiv \bar{b}_s - \bar{b}_w$). We use Mann-Whitney two-sample test (MW) to assess the statistical significance of the difference. In the symmetric stage the difference is negative and not significant (p>0.1, MW, 2-tail). In both asymmetric stages the difference is positive and statistically significant at the 1 percent significance level (MW, 1-tail). This finding is consistent with the prediction that the strong bidders' bid distribution stochastically dominates that of weak bidders. However, this is not a particularly strong test of the theory. Note that the strong bidder's value distribution first-order stochastically dominates that of the weak bidder: $F_s(v) < F_w(v)\ \forall v \in (0, \bar{v})$. This implies that $E_{F_s}[\beta(v)] > E_{F_w}[\beta(v)]$ for any monotone function $\beta(.)$ (Hadar and Russell (1969)). In other words, the average strong bid is expected to be higher than the average weak bid even if the strong and the weak bidders use the same bidding strategy.

In contrast, the predicted pointwise relation between bids, $\beta_s(v) < \beta_w(v)\ \forall v \in (0, \bar{v})$, is a stronger test of the theory as it goes against the ranking of mean bids. The hypothesis is also more difficult to evaluate as the theory does not provide a parametric model to efficiently use the data. The approach we take is to compare average bids for each of the eight intervals of values used in Table 2. Statistically significant negative differences in average bids $(\bar{b}_s - \bar{b}_w)$ can be taken as evidence in support of the hypothesis. This conclusion does not depend on our particular choice of intervals for the following reason. Given the value distribution functions in our setup, if we truncate the range of values to some interval $[x, y] \subset (0, \bar{v})$ then for the truncated value distribution functions on the interval, $\bar{F}_s(v) \equiv \frac{F_s(v) - F_s(x)}{F_s(y) - F_s(x)}$ and $\bar{F}_w(v) \equiv \frac{F_w(v) - F_w(x)}{F_w(y) - F_w(x)}$, the stochastic dominance is preserved: $\bar{F}_s(v) < \bar{F}_w(v)\ \forall v \in (x, y)$.[14] Thus, for *any* interval $[x, y]$ it continues to hold that $E_{\bar{F}_s}[\beta(v)] > E_{\bar{F}_w}[\beta(v)]$ for any monotone $\beta(.)$. Consequently, if there is no difference in behavior between strong and weak bidders, the average bid of the strong bidders on the interval is expected to be higher than the average weak bid. It can only be the case that $E_{\bar{F}_s}[\beta_s(v)] < E_{\bar{F}_w}[\beta_w(v)]$ if $\beta_s(v)$ is sufficiently below $\beta_w(v)$ on the interval. Therefore, if we do find that that average strong bid is lower than the average weak bid $(\bar{b}_s < \bar{b}_w)$ it should serve as convincing evidence of strong bidders bidding systematically lower than weak bidders for that interval of values. Note also that the condition $E_{\bar{F}_s}[\beta_s(v)] < E_{\bar{F}_w}[\beta_w(v)]$ may not hold on a large interval even if $\beta_s(v) < \beta_w(v)$ holds. Recall that $E_{\bar{F}_s}[\beta_s(v)] > E_{\bar{F}_w}[\beta_w(v)]$ when the interval includes all possible values. Therefore, the intervals we choose should be sufficiently small so that the stochastic dominance of the strong bidders' value distribution does not outweigh the pointwise relationship between the bidding function. At the same time, the intervals can't be too small as the number of observations would be insufficient and the power of a statistical test would be too low. The eight intervals in Table 2 were chosen with this trade-off in mind.

From Table 2 we can conclude that in both stage 2 and stage 3 average bids are higher for weak bidders on high value intervals. The result is statistically significant on intervals 3 and 5-7 (p<0.1, MW, one-tail). In stage 3 the difference tends to be higher in magnitude and statistically significant on interval 0 as well (p<0.1, MW, one-tail). These results are consistent with the Nash equilibrium (NE) predictions regarding asymmetry in bidding behavior. In the NE strong bidders bid lower than weak bidders regardless of the level of risk aversion, as can be seen from Table 3. Pezanis-Christou (2002) and Guth et al. (2005) also find that weak bidders tend to submit higher bids than strong bidders when having the same value for the item.

Observation 2. *Consistent with the NE, weak bidders tend to bid higher than strong bidders in stages 2 and 3.*

In the symmetric stage we should not observe any difference in bidding behavior between the bidder categories since all bidders are symmetric. However, our results indicate that the strong bidders tend to bid lower than the weak bidders in stage 1 even though their values are drawn from the same distribution. We find that the difference in average bids is not significant (p>0.1, MW, 2-tail). This asymmetry in bidding could potentially be attributed to

[14] Using the value distribution functions of the bidders: $\bar{F}_s(v) - \bar{F}_w(v) = \frac{F_s(v) - F_s(x)}{F_s(y) - F_s(x)} - \frac{F_w(v) - F_w(x)}{F_w(y) - F_w(x)} = \frac{v^2 - x^2}{y^2 - x^2} - \frac{v - x}{y - x} = \frac{(v-x)[v-y]}{(y-x)(y+x)} < 0$.

the within-subject design of our experiment. Subjects may be carrying over their experience in the asymmetric stages to the symmetric stage. For strong subjects in particular the experience in the asymmetric stages may teach them that higher profits can be earned by bidding lower even in the symmetric stage. We re-visit the issue of foregone profits when performing best reply analysis.

4.1.3. Bidding over time and across sessions

The asymmetry in stage 1 also hints at potentially important learning trends in subjects' behavior. To explore this idea we construct time series of the ratio of observed bids (b_i) to $\text{RANE}_{0.5}$ bids (β_i).[15] The time series are shown in Figure 7 in the appendix. Six panels contain data for strong and weak bidders in each of the three stages. Each panel in the figure depicts three time series: one for each of the three sessions conducted. The series from the same session have the same line pattern across panels. Period averages of the ratio b_i/β_i are plotted. The series are smoothed using a 3-period moving average. The time series illustrate how the observed bids compare to the $\text{RANE}_{0.5}$ bids over time. When the two coincide, the ratio is equal to one (a solid line at one is provided as a benchmark). Downward patterns could be indicative of learning. Inspecting the figure we note several things. First, time patterns vary across bidder categories, sessions, and stages. For example, in session 2 there is a clear downward trend for weak bidders while no trend is discernible in session 3. This variety suggests that focusing on a particular subset of data under an assumption about learning trends would be misleading. As a result, we report our results based on all the available data. Second, there is substantial heterogeneity across sessions in the level of bidding. It serves as an illustration why a between-subject version of this experiment would be particularly inefficient. Differences between sessions are affected by differences in subjects' characteristics (such as risk aversion) as well time paths that could be influenced by a number of uncontrollable factors. Such heterogeneity would mask the treatment effect and require a large number of sessions for each treatment condition to detect a statistically significant difference (if any). Finally, Figure 7 illustrates Observation 1. While $r_i = 0.5$ is a good descriptor of the average level of bidding in stage 1, in the asymmetric stages strong bidders clearly bid higher than the NE with the same level of risk aversion.

4.1.4. Changes in bidding between stages 1 and 2

Bidding changes between stages are less consistent with the NE than differences in bidding between bidder categories. Strong bidders do seem to decrease their bids between stage 1 and stage 2 on all value intervals except for 0 and 1 (see Table 2). However, the decrease is quite small. On the other hand, weak bidders seem to *increase* their bid slightly for some intervals of values contrary to the NE predictions. Our design allows us to perform within-subject tests to assess the statistical significance of the change in behavior between the stages. We perform tests on two measures. The first one looks at the percentage of subjects who decrease their bids between stages more frequently than increase them. A dummy variable is constructed which

[15] Looking at time patterns of absolute bids would be misleading as it would also reflect any random pattern in the realization of values. We chose to normalize the observed bids using $\text{RANE}_{0.5}$ bids (functions of values) for an easy comparison between the two.

Table 4: Bid changes between stages

		Between stages 1 and 2		Between stages 2 and 3	
		Strong	Weak	Strong	Weak
Values >250	Decrease bids more often	73%*** (29/40)	30%* (6/20)	50% (20/40)	30%* (6/20)
	Difference in average bids	$-11.4^{\sharp}$	3	-2.1	$6.1^{\ddagger}$
All values	Decrease bids more often	73%*** (29/40)	40% (8/20)	45% (18/40)	35% (7/20)
	Difference in average bids	$-11.1^{\sharp}$	2.2	-0.8	$4.2^{\ddagger}$

* - p<0.1, ** - p<0.05, *** - p<0.01; Binomial test, one-tail, Null: percentage = 50%;

†- p<0.1, ‡- p<0.05, $^{\sharp}$- p<0.01; Wilcoxon matched-pairs signed-rank test;

takes on the value of one if a subject decreases his or her bids more often than increases.[16] We perform a one-tail Binomial test (BIN) on this variable under the null hypothesis that there is no systematic tendency for subjects to change their bids in a particular way, i.e. roughly 50 percent tend to decrease their bids while the other 50 percent tend to increase them. The results are reported in the first row of Table 4 ("Decrease bids more often"). For example, the first cell of the table tells us that between stage 1 and stage 2 73 percent of the strong bidders (29/40) decrease their bids more often than increase. This percentage is statistically different from 50 percent (p<0.01, BIN, one-tail). While this is a measure of frequency with which bidders decrease their bids, the other measure is aimed at assessing the magnitude of bid changes. The second row of Table 4 reports the change in average bids between any two stages ("Difference in average bids"). Thus, the first cell in the second row reports that for the strong bidders the change in the average bid between stage 1 and stage 2 is -11.4. Statistical significance of the change is assessed using the WILC test performed on subject averages in the two stages. Apart from the whole sample Table 4 also includes the results for middle to high values only. It can be argued that at lower values subjects may have little incentive to bid optimally as the probability of winning and the potential payoff are low. We used an arbitrary cut-off point of 250 to see if our conclusions change when we focus on a region where stakes are relatively high.

The results shown in Table 4 are consistent with the conclusions derived from the analysis of average bids. Strong bidders have a tendency to decrease their bids when they enter stage 2. This conclusion is supported by both tests. Thus, 73 percent (29/40) of strong bidders decrease their bids more often than they increase them between stage 1 and stage 2 (p<0.01, BIN, one-tail). The average decrease in bids is -11.4 (p<0.01, WILC, two-tail). Although the magnitude of the decrease is much smaller than predicted by the Nash equilibrium, it is statistically significant. The results are unchanged if we restrict our attention to the middle-high values. On the contrary, weak bidders show no tendency to decrease their bids as they

[16] For weak bidders the application of the test is straightforward. For strong bidders we can only use observations where a subject's signal in Stage 1 becomes the team's value in asymmetric stages. These are the only observations to which the comparative static results are applicable.

should according to the Nash equilibrium. Instead, only 40 percent (8/20) of weak bidders decrease their bids more often than they increase them. At middle-high values the percentage is even smaller. In fact, at the 10 percent significance level we can assert that weak bidders tend to increase their bids between stage 1 and stage 2 when their values are high (70 percent or 14/20 increase more often, $p<0.1$, BIN, one-tail). Quantitatively, the average change in bids is positive for weak bidders, but not statistically different from zero (2.2 for all values and 3 for middle-high values).

Observation 3. *Consistent with the NE, strong bidders tend to decrease their bids between stage 1 and stage 2. Contrary to the NE, weak bidders are more likely to increase their bids between stage 1 and stage 2 than to decrease them.*

4.1.5. Changes in bidding between stages 2 and 3

The only difference between stages 2 and 3 is that in stage 3 strong bidders belonging to the same team can communicate with each other before submitting a bid. If individual strong bidders and strong teams are identical in their preferences one would expect no change in behavior between stages 2 and 3. A team of two bidders would submit the same bid as each of them would on their own. One possible reason for a change is if the team has different risk preferences compared to its members. The relevant concept in psychology is risky shift - a tendency for groups to become less risk-averse than individuals. The opposite tendency is referred to as cautious shift. The left panel of Figure 2 shows that if the strong team becomes less risk-averse and this fact is common knowledge then both strong and weak bidders should decrease their bids pointwise between stages 2 and 3. On the other hand, if a cautious shift occurs, the bidders should increase their bids. Another factor that may account for a change in bidding is the principle that "two heads are better than one". As suggested by Observation 1, Nash equilibrium with risk aversion roughly calibrated to fit the symmetric stage does not explain well the behavior of strong bidders in the asymmetric stages. They bid higher compared to the NE bidding function. Thus, strong bidders may be able to improve their expected utility by submitting lower bids. Then, a change in the direction of lower bids on the part of strong teams can be interpreted as better decision making.

As Table 4 shows, strong bidders are not affected by the change in the bidding format. Strong bidders appear to be roughly equally divided as to whether they tend to increase their bids or decrease when going from stage 2 to stage 3. On the other hand, weak bidders seem to respond to the ability of the strong team members to communicate. The effect is small but statistically significant and is stronger at middle-high values. The change between average weak bids is positive regardless of the sample considered (4.2 for all values and 6.1 for middle-high values, $p<0.05$, WILC, 2-tail). Additionally, if we focus on values greater than 250, only 30 percent (6/20) of weak bidders decrease their bids more often than increase ($p<0.1$, BIN, one-tail). In other words, 70 percent of weak bidders tend to increase their bids between stage 2 and stage 3 when their values fall into the middle-high region. Thus, weak bidders seem to bid slightly higher when facing a 2-member team even though the value of the object to the team does not change. This result is partially consistent with the cautious shift hypothesis which would induce higher bidding on the part of both strong and weak bidders. A risky shift would induce lower bidding while two-heads-better-than-one phenomenon would not affect the behavior of weak bidders. However, cautious shift would directly affect strong bidders who do not seem to change their behavior. A possible explanation that does not rule out a cautious shift is that the strong bidders' tendency to bid higher due to a higher level of risk aversion is counteracted by an incentive to improve their utility by bidding lower originating from superior decision making abilities of the group. On balance, Weak bidders may be bidding higher anticipating a cautious shift from 2-member teams that does not actually materialize.

Observation 4. *Group decision-making in stage 3 leaves the bidding behavior of strong bidders roughly unchanged, while weak bidders tend to increase their bids.*

On the surface, group decision making has little effect on strong bidders' behavior in our data. One potential explanation is that members of a team simply compromise if their stage 2 bids happen to be different. If they simply took the average of the two bids as the team's bid, the observed result would obtain. We take a deeper look into the decision making process to test this hypothesis. First, we note that an agreement is reached between team members about 87 percent of the time. In the remaining cases the team members submit different stage 3 bids, one of which is randomly chosen to be the team bid. Next, we compute the frequency with which each stage 3 bid falls within the interval formed by the two team members' stage 2 bids. It turns out that only 65.2 percent of stage 3 bids are contained within the corresponding (closed) intervals and can be considered as a compromise in reference to the stage 2 bids. Interestingly, 14.0 percent of bids formed after the communication phase are above both stage 2 bids, while 20.8 percent are below. Thus, the group decision making process goes beyond a simple compromise. We also regress stage 3 bids on the maximum and the minimum of the two team members' bids using OLS (with robust standard errors allowing for subject-based clusters) and random effects models. Both specifications produce similar results significant at better than 5 percent level. The higher team member's bid has approximately twice the influence in determining stage 3 bid: the coefficients are 0.614 (OLS) and 0.623 (RE) for the higher bid compared to the coefficients of 0.305 (OLS) and 0.268 (RE) for the lower bid.

In an attempt to gain insight into these patterns we examined the messages exchanged between strong team members during the communication phase. Strong bidders made occasional references to risk in sentences like "better safe than sorry" and "should we risk it this time?". Subjects also discussed the history and the patterns of play by their weak opponents. However, we were not able to detect any systematic changes in bidding following such messages.

4.1.6. Bidding function estimates

One way to interpret our results so far is that the two bidder categories respond to value distribution asymmetry due to joint bidding by increasing the asymmetry in their bidding behavior. Weak bidders tend to bid higher than strong bidders. However, in the Nash equilibrium a decrease in the number of bidders following the creation of the joint-bidding entity also allows both bidder categories to substantially lower their bids. Subjects do not seem to respond appropriately to this strategic force. Figure 4 explores this and other features of the observed bidding behavior using estimated bidding functions. We estimate bidding functions by fitting a quadratic polynomial to the pooled data from all the stages.[17] Dummies for the three stages and all the interaction terms are included. Table 10 (see Appendix)

[17] Higher order terms were not statistically significant.

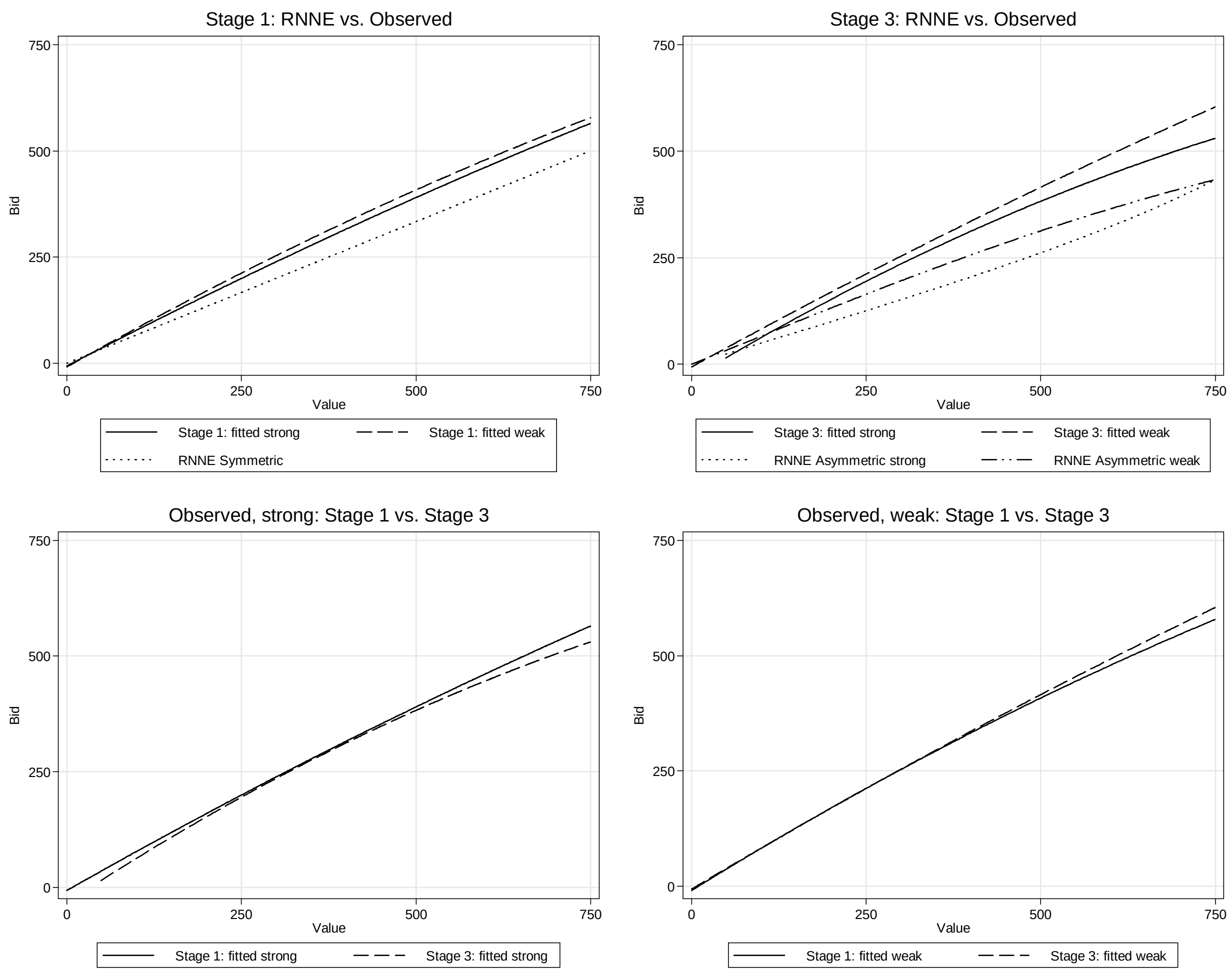

Figure 4: Estimated bidding functions (random effects regression)

contains the estimation results from a random effects model and an OLS model with robust standard errors allowing for clusters formed by observations from the same subject. Figure 4 plots the estimated bidding functions from the random effects model. Focusing on the overall effect of joint bidding, the lower two panels compare the estimated bidding functions in stage 1 and stage 3 for each bidder category. The lower-left panel shows that in the asymmetric stage strong bidders bid lower mostly at high values consistent with the NE prediction. The lower-right panel, however, shows that weak bidders tend to increase their bids, contrary to the theoretical predictions. The top two panels illustrate the differences in bidding between the two bidder categories as well as between the estimated and the RNNE bidding functions. The $\text{RANE}_{0.5}$ bidding functions share key features of the RNNE bidding functions but are pointwise higher (not depicted). The upper-left panel shows the estimated bidding functions from stage 1 and the RNNE in the symmetric environment. It is clear that subjects overbid relative to the RNNE. A small asymmetry between the two bidder categories is also present. In the upper-right panel, estimation results are plotted for stage 3 together

Table 5: Seller's revenue

	Stage 1	Stage 2	Diff12	Stage 3	Diff23
Sample average	449.81	440.92	$-8.89^{*\sharp}$	441.47	0.55^{*}
(sd)	(134.54)	(132.97)		(135.72)	
Re-matched	448.22	440.78	-7.44	440.75	~ 0
(sd)	(2.469)	(4.097)		(2.51)	

* - p<0.1; Binomial test, one-tail, Null: equal number of negative and positive changes between stages;

$\sharp$- p<0.01; Wilcoxon matched-pairs signed-rank test;

with the RNNE predictions for the asymmetric environment. The panel indicates that the gap between the estimated bidding functions widens in stage 3 as the two bidder categories tend to bid differently at high values. The panel emphasizes that the nature of the observed asymmetry between the two bidder categories is different from the NE prediction. According to the latter, the asymmetry peaks in the middle of the value range and diminishes at higher values. This is not the case with the observed bidding behavior. Additionally, the panel shows that substantial overbidding relative to the RNNE is present in the asymmetric environment as well.

4.2. Seller's revenue

Based on the theoretical predictions (see Table 1) we expect the seller's revenue to drop between the symmetric and the asymmetric markets by 15 percent when bidders are risk-neutral. If bidders are CRRA risk-averse with $r_i = 0.5$, the expected decline is smaller at about 10 percent. Table 5 reports the average seller's revenue observed in the experimental data. While the observed revenue is much higher than the RNNE prediction of 375, it is almost exactly equal to the $\text{RANE}_{0.5}$ revenue of 450. Consistent with the Nash equilibrium, the seller's revenue falls between stage 1 and stage 2, although the change is much less pronounced than predicted by theory. On average, the revenue seems to decline by only about 2 percent between stage 1 and 2 ("Diff12" column) which is much less than the 10 percent decrease in the $\text{RANE}_{0.5}$ case. Such a small decline is not surprising in light of our findings on bidding behavior. In the Nash equilibrium both bidder categories should decrease their bids when moving from the symmetric to the asymmetric environment. However, the subjects in the strong role do not decrease their bids nearly as much as predicted while the subjects in the weak role tend to increase their bids. The difference in the observed revenues between stage 1 and stage 2 is so small that a between-group equality of means test does not yield any significant results (p>0.1, MW). However, within-group tests pick up the decline (p<0.01, WILC; p<0.1, BIN).[18] A comparison of the asymmetric stages with and without team communication yields no conclusive results ("Diff23" column). The difference is positive and borderline significant using a one-tail binomial test (p<0.1). It indicates that the revenue increases more frequently than it decreases between stage 2 and stage 3. This result would be consistent with the observation that weak bidders bid higher in stage 3. However, the magnitude of the effect is negligible.

[18] Within-group tests compare pairs of outcomes in different stages by matching the outcomes produced by the same group. Between-group tests compare two samples of outcomes in different stages without taking advantage of this information.

Observation 5. *In stage 1 the revenue is as predicted by the RANE$_{0.5}$. Consistent with the Nash equilibrium, revenue decreases between stage 1 and stage 2. However, the decrease is much smaller than predicted. Group decision-making has almost no effect on the revenue.*

The observed effect on revenue is thus very small. At the same time, in small samples the matching of subjects into groups may have a sizable effect on revenue. The standard procedure of random group assignment of subjects during experiments adds an additional layer of noise that makes simple averages of market-level measures less reliable in small samples. This is especially true for auctions because swapping bidders in any two markets may have a profound effect on allocations and prices. To verify that our results are not an artifact of the random group assignment we perform simulations by randomly re-matching subjects 1000 times and report the averages obtained from these simulations. They constitute a more precise estimate of the market-level implications of the observed bidding behavior. In the second row of Table 5, the distributions of means of 480-observation samples are reported (20 groups $\times$ 24 periods, see Table 11 in the Appendix). The distributions are statistically indistinguishable from Gaussian (skewness and kurtosis test for normality). Overall, these results support the conclusions based on the observed average revenues. Revenue tends to be lower in the asymmetric environment by about 2 percent, while communication between team members seems to have no effect on the revenue.

4.3. Allocative efficiency

Inefficiency in auction markets is a concern for policy makers. In symmetric auctions inefficiency is absent as the ranking of bidders' values translates directly into the ranking of submitted bids. As a result, the bidder with the highest value always wins. In contrast, in auctions with ex-ante asymmetry in value distributions, the Nash equilibrium strategies of weak and strong bidders differ. As a result, a weak bidder with the lower value wins the auction with a positive probability. In the model studied here the drop in efficiency when moving from the symmetric to the asymmetric environment is around 9 percent if both bidders are risk-neutral. In the RANE$_{0.5}$ the decline is 7 percent. We analyze the efficiency implied by the data by performing random re-matching as described above. It is especially important in this context since perturbations in group assignment have significant effect on allocative efficiency, which is a binary variable.

Inefficiency in symmetric auctions has been observed in other experimental studies (see Kagel (1995)). It can be attributed to random errors and inconsistencies in bidding behavior. The results here are no exception. As reported in Table 6, inefficiency is substantial in stage 1. In more than 14 percent of cases the object goes to the bidder who does not have the highest value. This result is illustrated in Figure 5. In the figure the set of possible value 3-tuples (a value draw for each symmetric bidder) is represented by an interval. The interval is divided into three regions depending on which bidder's value draw is the highest of the three. The outcomes associated with the Nash equilibrium are shown in the lower part of the figure. If all bidders follow the Nash equilibrium strategy each bidder wins whenever his value is the highest. The upper part of Figure 5 shows how the set of value 3-tuples is partitioned by the experimental data. The possible areas are labeled 1a through 3c. The sum of areas 1b, 1c, 2a, 2c, 3a, and 3b corresponds to the 14 percent observed inefficiency in the symmetric stage. These areas are absent in the Nash equilibrium.

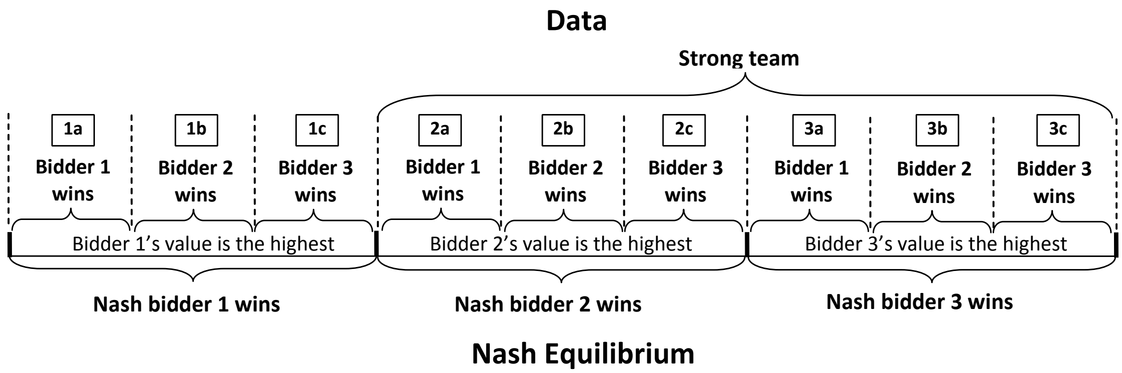

Figure 5: (In)efficient allocations in the symmetric setup: the Nash equilibrium vs. data

Furthermore, contrary to the theoretical prediction that efficiency should decline with joint bidding, our data suggest otherwise. Efficiency rises in asymmetric markets by about 2 percentage points. One plausible explanation is that in high-noise environments efficiency is responsive to the number of bidders. Specifically, the average distance between the highest and the second-highest values decreases as the number of bidders in the market increases. Consequently, other things being equal, if bidding has a random error, inefficient allocations are more likely when more bidders participate in the auction. This force may dominate the strategic adjustment, which is responsible for higher inefficiency in a NE of the asymmetric auction. To measure the effect of a smaller number of bidders we calculate the "intra-strong" inefficiency: those inefficient allocations in stage 1 that are caused by strong bidders winning the auction when their team partners should have won it instead. This type of inefficiency is only possible in the symmetric stage where strong bidders from the same team compete against each other. Figure 5 clarifies the origin of intra-strong inefficiency in the symmetric environment. Suppose bidders 2 and 3 are in the same team so that they bid jointly in asymmetric stages. Areas 2c and 3b represent the outcomes in which a team member wins when the other team member has the highest value in the symmetric stage. As reported in 6, this type of misallocation occurs in 3.5 percent of auctions and, by definition, disappears in the asymmetric stages where those bidders become a team. Note that inefficiency falls by at most 2 percentage points between the symmetric and the asymmetric stages which suggest that other types of inefficiency actually rise. On balance, the noise-reducing impact of fewer bidders outweighs the opposing forces.

Observation 6. *Contrary to the Nash equilibrium, efficiency is higher in the asymmetric environment. The increase in efficiency between the symmetric stage and the asymmetric stages is associated with a decrease in the number of bidders.*

To assess the strength of the force that pushes efficiency up in the asymmetric stages we perform simulations along the lines of Gode and Sunder (1993) using "zero-intelligence" bidders. Zero-intelligence bidders submit random bids with the only restriction that their bids are below their values. With such bidders we estimate the efficiency to be about 63 percent in the 3-bidder symmetric auction (400,000 auctions). In this case the expected distance between

Table 6: Allocative efficiency

	Stage 1	Stage 2	Stage 3
Inefficiency	0.142	0.126	0.122
(sd)	(.013)	(.014)	(.013)
Intra-strong	0.035	0	0
(sd)	(.005)		

the first- and second-highest values is 1/4 of the value range.[19] As we move to the asymmetric environment the expected distance rises to 1/3 of the range. As a result, the efficiency rises to 72 percent.[20] Thus, the change in the auction structure has a powerful positive effect on efficiency when bidding is completely random. Subject behavior is more deterministic and therefore the effect is not as strong but still observable.

4.4. Bidders' payoffs and a best-response analysis

Waehrer (1999) showed that in the equilibrium of this model larger groups have a smaller per member payoff than smaller groups. Thus, joint bidding has a positive externality on the bidders who are not part of it. If the decision to bid jointly were endogenous, this would create a free rider problem, whereby the incentives to bid jointly would be weakened even though this change in auction structure is profitable for all the participants. While in our experiment the decision to bid jointly is exogenously imposed on the bidders we can nonetheless assess the ex-ante incentives by comparing the observed payoffs before and after the change.

The Nash equilibrium prediction is that in stages 2 and 3 the payoffs of all bidders are expected to be higher than in stage 1. When bidders are risk-neutral the predicted increase is smaller for strong bidders than for weak bidders in line with the results in Waehrer (1999). The per-member profit of strong bidders is expected to increase by about 20 percent, while the profit of weak bidders is expected to rise by more than 40 percent. This prediction does not hold in the data. Tables 7 and 8 (for strong and weak bidders respectively) report profit conditional on winning, probability of winning and average profit obtained from re-matching simulations (the "Observed" portion of the tables). As predicted by the RNNE, the decrease in the number of bidders due to joint bidding is profitable for both strong and weak bidders. Per-member average profit increases for both bidder categories by about the same amount. In percentage terms, between stage 1 and stage 2, strong bidders enjoy an approximately 7.3 percent boost in average per-member earnings (34.19 to 36.69, Table 7) while weak bidders' profits increase by 6 percent (30.75 to 32.61, Table 8). Additionally, between stage 2 and stage

[19]For the sample of size n drawn from the uniform distribution on $[0, 1]$, the difference between any two adjacent order statistics has the density $n(1-x)^{n-1}$. Thus, for $n = 3$ the expected value of the difference is 1/4, while for $n = 2$ the expected value is 1/3. Curiously, for $n = 2$ where one observation is uniform while the other is the maximum of two draws from the uniform distribution (our asymmetric case), the density of the distance is the same, so that the expected value is again 1/3.

[20]Moving from the symmetric 3-bidder to the symmetric 2-bidder case the efficiency rises from 63% to 75%. The subsequent fall in efficiency to 72% when one of the bidders is advantaged is not very intuitive. It can be attributed to the fact that the expected maximum value increases in the 2-bidder asymmetric case compared to the 2-bidder symmetric case extending the upper bound of the interval from which bids are drawn. As a result, the dispersion of random bids increases, contributing to inefficiency. This force is much smaller than the effect of the number of bidders.

Table 7: Strong bidders' payoffs

	Observed				Best response	
	Cond.	Prob. of	Per member		Per member	
	Profit	winning	Profit	RA Utility	RN	RA
Stage 1	105.9	0.645	34.19	6.64	49	7.46
(sd)	(2.91)	(.012)	(1.154)			
Stage 2	119.1	0.615	36.69	9.49	62.52	10.98
(sd)	(4.66)	(.015)	(1.675)			
Stage 3	123.5	0.607	37.62	9.54	61.8	10.8
(sd)	(3.43)	(.013)	(1.338)			

3, the profit of strong bidders increases further by about 2.5 percent while that of the weak bidders drops by about 2.8 percent. These changes are small compared to those expected in the RNNE. Furthermore, asymmetric auctions seem to be more profitable for strong bidders or at least weak bidders do not reap unusually high benefits compared to strong bidders.

Observation 7. *Contrary to the RNNE prediction, weak bidders do not obtain a higher increase in expected profits than strong bidders as a result of joint bidding.*

The above pattern of changes in profits is more compatible with the $\text{RANE}_{0.5}$. Although the predicted increases in expected profits for both bidder roles are still higher than observed (the former is about 38 percent for both, Table 1), the strong bidders are expected to gain almost as much as the weak bidder, similar to what we find in the data. Additionally, note that when bidders are risk-averse they maximize expected utility rather than profits, therefore, it is more appropriate to look at the former. In expected utility terms, the prediction that strong bidders do not gain as much as weak bidders is completely reversed. Strong bidders are predicted to increase their expected utility by 57 percent while weak bidders benefit by a more modest 29 percent (Table 1). In "Observed $\rightarrow$ Per member $\rightarrow$ $\text{RA}_{0.5}$ Utility" columns of 7 and 8 we report the "observed" expected utilities under the assumption that subjects are risk-averse with $r_i = 0.5$. These utilities are calculated from the observed probabilities of winning and conditional profits reported in the same tables. Based on these numbers we find that subjects in the role of strong bidders actually increase their expected utility by $45 - 47$ percent between the symmetric and the asymmetric stages while weak bidders increase their utility by only 7 percent. There are a couple of things to note about these results. First, they suggest that the Nash equilibrium with risk aversion is at least partially able to explain the observed discrepancy in expected gains that favor strong bidders. Therefore, if joint bidding were an endogenous choice, risk aversion should be viewed as an important bidder characteristic that affects the incentives to bid jointly. Second, the $\text{RANE}_{0.5}$ may actually underestimate the degree to which strong bidders gain more relative to weak bidders. Thus, bidders may have a bigger incentive to be part of a joint-bidding arrangement than suggested by the RNNE or even the $\text{RANE}_{0.5}$.

Observation 8. *Under the assumption of risk aversion, strong bidders enjoy a much higher increase in expected utility as a result of joint bidding than do weak bidders. The relative gain of strong bidders is higher than predicted by the* $RANE_{0.5}$.

Table 8: Weak bidders' payoffs

	Observed				Best Response	
	Cond. Profit	Prob. of winning	Per member		Per member	
			Profit	$RA_{0.5}$ Utility	RN	$RA_{0.5}$
Stage 1	85.89	0.355	30.75	6.58	52.22	7.85
(sd)	(4.18)	(.012)	(1.836)			
Stage 2	83.74	0.385	32.61	7.05	57.01	8.30
(sd)	(3.6)	(.015)	(1.835)			
Stage 3	79.8	0.393	31.7	7.02	56.84	8.39
(sd)	(3.23)	(.013)	(1.6)			

Thus, we find that in relative terms the gains of weak and strong bidders as a result of join-bidding are partially consistent with the Nash equilibrium with risk aversion. However, quantitatively there are significant discrepancies. Thus, the observed bidders' payoffs are lower than those in the $RANE_{0.5}$ and so are the changes in payoffs due to joint bidding. A natural question arises: can subjects actually improve their expected utility or do they approximately best-respond given the behavior of their opponents? To answer this question we construct (exact) best response functions for both bidder categories in every stage. For every value we find the bid that brings the highest expected utility where the probability of winning is based on the opponent's bid distribution in the data. Figure 6 plots these functions for risk-averse bidders with $r_i = 0.5$ ($RABR_{0.5}$).[21] In the upper-left panel, stage 1 $RABR_{0.5}$ functions for both bidder categories are compared to the estimated bidding functions in the symmetric environment. In the upper-right panel, stage 3 $RABR_{0.5}$ for both bidder categories are compared to the estimated bidding functions in the asymmetric environment.[22] The lower two panels compare the $RABR_{0.5}$ functions in stage 1 and stage 3 for strong (left panel) and weak (right panel) bidders. The upper-left panel shows that in terms of bid magnitude, the $RABR_{0.5}$ functions track the estimated bidding functions rather well in stage 1. In stage 3, however, the RABR functions seem to be substantially lower than the estimated bidding functions. This result is consistent with our earlier observation that strong subjects do not decrease their bids in asymmetric stages as much as predicted by the $RANE_{0.5}$. In addition, the $RABR_{0.5}$ functions suggest that weak bidders too could benefit from reducing their bids in the asymmetric stages.

To understand why the bidders don't reduce their bids we compute the expected utilities that risk-averse bidders could have obtained had they followed the $RABR_{0.5}$ bidding functions.[23] The utilities are reported in Tables 7 and 8 ("Best response → Per member → $RA_{0.5}$" columns). When we compare them to the actual expected utilities that bidders obtained under the assumption of $r_i = 0.5$, it turns out that strong bidders could have improved their

[21] The exact best response is a step function because the sample of the opponent's bids is finite and the probability of winning is itself a step function.

[22] Best response functions in stage 2 are very similar to the ones in stage 3

[23] A best-response expected utility is obtained by numerical integration. It is calculated as a weighted average of the best-response expected utilities associated with every possible value draw. The weights are the associated probabilities of drawing a certain value ($\frac{1}{750^2}v$ for the strong bidders in the asymmetric stages, and $\frac{1}{750}$ for all the other cases).

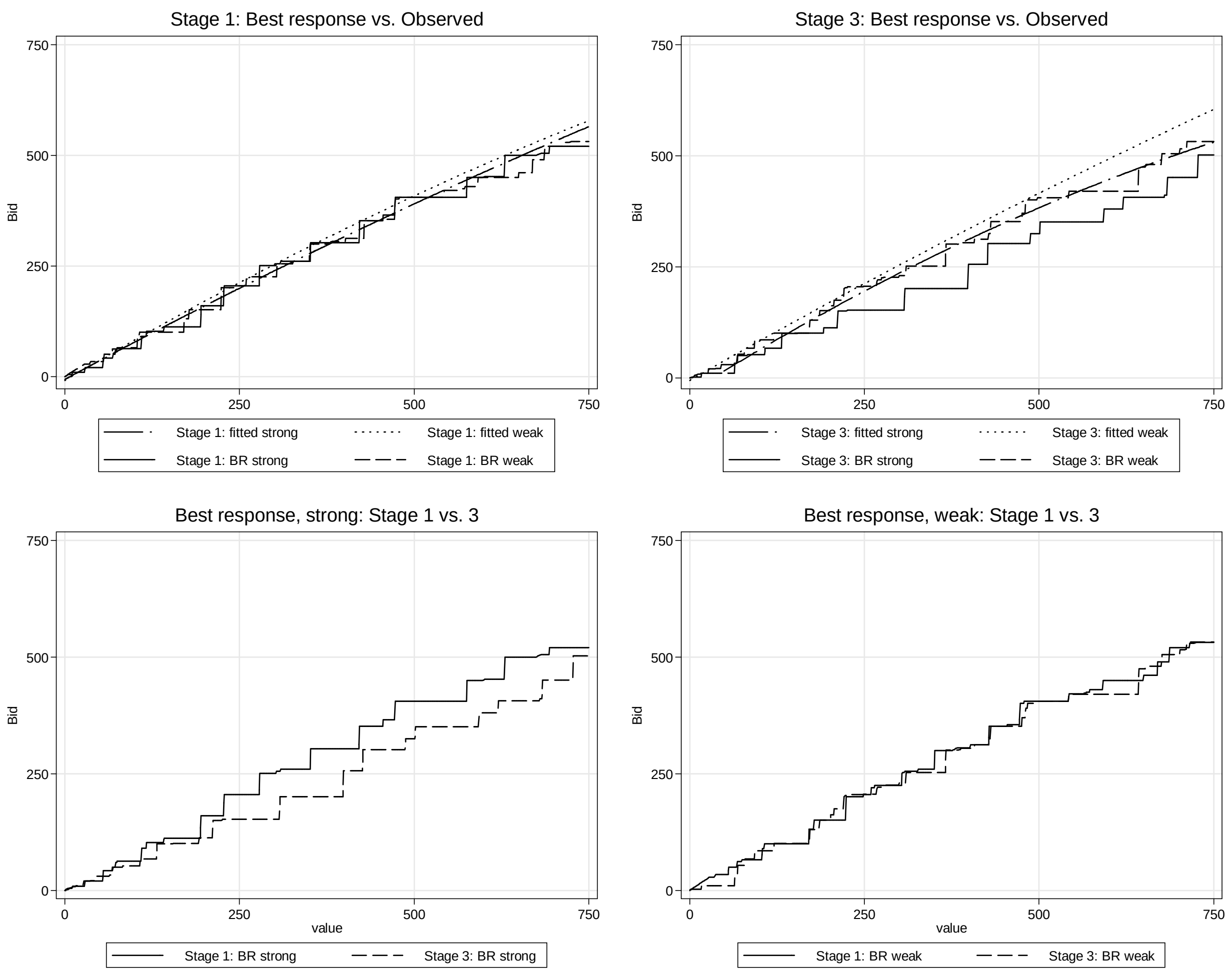

Figure 6: Best response by risk-averse bidders ($r_i = 0.5$) and estimated bidding functions.

utility by about the same percentage in all stages (12, 16, and 13 percent in respectively stages 1, 2, and 3). Thus, the incentives to use the $\text{RABR}_{0.5}$ strategies were about the same across stages. In other words, subjects came as close to the $\text{RABR}_{0.5}$ in stage 1 as in stage 3 in terms of forgone expected utilities. Similarly, the weak bidders could have improved their utility by a fairly narrow range of 18-20 percent in all stages. Thus, what looks like substantially sub-optimal bidding in stage 3 turns out to be no more sub-optimal than that in stage 1. It seems that in the asymmetric stages larger deviations from the $\text{RANE}_{0.5}$ fall within a similar margin of error (in expected utility terms) as in stage 1.[24]

[24]For comparison we perform a similar exercise assuming that subjects are risk neutral (risk neutral best response functions or RNBR, not depicted). We find that RNBR functions are in the neighborhood of the RNNE functions (see Figure 4, which compares estimated bidding functions to the RNNE). It is interesting that despite subjects' overbidding relative to the RNNE, the RNBR functions are not that different from RNNE in terms of bid magnitude. Thus, scale-wise best response functions are quite insensitive to overbidding by

Observation 9. *Under the assumption that subjects are risk-averse, we find that for strong bidders expected utility losses relative to best responses are similar across stages. The same holds for weak bidders.*

A related question is how gains from joint bidding that we observe in the data compare to those obtainable by following the $\text{RABR}_{0.5}$. The lower-right panel of Figure 6 shows that the $\text{RABR}_{0.5}$ function for weak bidders in stage 3 is hard to distinguish from that in stage 1. The associated $\text{RABR}_{0.5}$ expected utility increases by 7 percent between stage 1 and stage 3 (see Table 8). Therefore, given the behavior of strong bidders there is relatively little additional incentive for weak bidders to lower their bids between the symmetric and the asymmetric stages as prescribed by the Nash equilibrium. This suggests that the predicted effect of joint bidding on weak bidders' expected utility is sensitive to the behavior of strong bidders. If strong bidders do not lower their bids to the extent predicted, the weak bidders' gains from joint bidding are noticeably reduced. This may be one of the reasons why we do not observe weak bidders submitting lower bids in the asymmetric stages. In fact, weak bidders obtain the same increase in expected utility (7 percent) by slightly increasing their bids. The lower-left panel indicates that strong bidders would find it beneficial to substantially lower their bids in asymmetric stages. The incentive to do so is significant: $\text{RABR}_{0.5}$ expected utility increases by 45 percent (see Table 7). Subjects in the role of strong bidders do reduce their bids as they move to the asymmetric environment, although the magnitude of the effect is lower than suggested by $\text{RABR}_{0.5}$. However, consistent with our discussion above, strong subjects manage to obtain largely the same increase in expected utility in spite of bidding higher than the $\text{RABR}_{0.5}$ (44 percent increase between stage 1 and stage 3).

Observation 10. *Subjects obtain expected utility gains from joint bidding similar to what they could obtain by following best response bidding functions.*

Putting these results together we conclude that the subjects' behavior can be viewed as an approximate best response consistent with the assumption of the CRRA risk aversion. With respect to the magnitude of bids, subjects experience similar utility loss across stages relative to the CRRA best responses. With respect to changes of bids between stages, they manage to boost their utility to a similar extent as they could have by following the best responses.

5. Conclusions

In this paper we explore an application of the asymmetric auction theory. We study ex-ante asymmetries that can result from such changes in the auction structure as legal joint-bidding ventures, collusive agreements, and mergers. In particular, we focus on the case where a group of symmetric bidders do not bid against each other, submit a single bid, and, in the case of winning, allocate the object to the member with the highest value. Such an agreement produces a party (strong) whose value distribution is more favorable compared to

one's opponents. With respect to profits, we find that by following the RNBR subjects could have improved their expected profits/utilities by 43-79 percent (see Tables 7 and 8). These are substantial incentives, and the subjects' failure to respond to them indicates that the assumption of risk-neutrality is unlikely to be appropriate.

the remaining bidders (weak). We obtain the Nash equilibrium bidding functions for a special case where two out of three symmetric bidders join to bid together. Both the equilibrium with risk-neutral bidders (RNNE) and the equilibrium with CRRA risk-averse bidders (RANE) are examined. Many features of the model have been characterized in the literature. Thus, in the RNNE of the resulting asymmetric market the weak bidder is predicted to bid pointwise higher than the strong bidder. We find that the result continues to hold in the RANE as long as the strong bidder is not much more risk-averse than the weak bidder. Of particular interest is the comparison between the outcomes before and after the joint-bidding arrangement is in place. All allocations are efficient in the original symmetric market, which is no longer the case in the asymmetric environment. In the latter case, the asymmetry in bidding functions induces equilibrium inefficiency since with positive probability the weak bidder wins the auction when his or her value is not the higher. Joint bidding also has an anti-competitive effect. Due to a reduction in the number of bidders, all bidders decrease their bids and the seller's revenue is reduced. On the flip side, profits increase for both the strong and the weak bidders. Under risk neutrality, weak bidders' profits are predicted to increase by more than the per-member profits of the strong entity. We find that this result is reversed if the bidders are sufficiently risk-averse.

We conduct a within-subject experiment designed to test the hypotheses outlined above. The results support the conclusion that subjects respond to the value distribution asymmetry due to joint bidding in a manner that is qualitatively consistent with the Nash equilibrium. Thus, weak bidders on average bid higher than strong bidders. Additionally, strong bidders tend to decrease their bids moving from the symmetric to the asymmetric environment as predicted. Quantitatively, subjects overbid relative to the RNNE. Risk aversion can account for higher bids. However, when the level of risk aversion is calibrated to fit the data in the symmetric environment, we find that strong bidders do not reduce their bids in the asymmetric environment as much as predicted by the RANE. Furthermore, contrary to the Nash equilibrium prediction weak bidders fail to adjust their bids downward with some tendency to move in the opposite direction. As a result, the decrease in the seller's revenue due to joint bidding is considerably smaller than predicted by either the RNNE or the RANE. We perform a best-response analysis and find that the observed behavior is not entirely inconsistent with some sort of risk-averse best-response. The comparative statics prediction with respect to efficiency changes is also violated as the efficiency increases in asymmetric markets. This deviation is consistent with a reduction in the number of bidders when bidding is noisy.

In addition, the experiment evaluates a behavioral aspect that can be associated with joint bidding: an increase in the number of decision-makers. Our model provides a prediction for the case when the effect is due to differences in risk aversion between individuals and groups. We find that the behavior of the strong entity is largely unaffected by an increase in the number of decision-makers, while weak bidders seem to respond by bidding higher. The RANE that takes into account group-individual differences in risk aversion cannot account for such behavior.

Finally, we find that contrary to the RNNE prediction joint bidding benefits strong bidders more than weak bidders. While the result is consistent with the RANE, we find that the theoretical prediction may underestimate the effect. The reason is that the strong bidders do not decrease their bids as much as predicted by the RANE. This type of overbidding does not hurt the strong bidders as much in expected utility terms, but it significantly reduces

	Firm 3: join	
	Firm 2	
	join	*not*
Firm 1 *join*	a, a, a	b, c, b
not	c, b, b	a, a, a

	Firm 3: not	
	Firm 2	
	join	*not*
Firm 1 *join*	b, b, c	a, a, a
not	a, a, a	a, a, a

Table 9: Endogenizing the decision to bid jointly

the ability of weak bidders to benefit from joint bidding. Thus, we find that the increase in the expected utility due to joint bidding is more than six times higher for the strong bidders compared to the weak bidders (vs. twice higher in the RANE). This result has important implications for the bidders' incentives to join forces. Using a simple model we can show that the relative gains of the two bidder categories play an important role. Consider three firms simultaneously deciding whether to merge or not. The normal form of the game is provided in Table 9. If only one firm chooses to merge the merger does not take place and all firms earn a. If two firms choose to merge, then the merger goes through, the merging firms earn $b > a$ while the remaining firm earns $c > a$. If all firms choose to merge, the merger does not occur, as a merger to monopoly is likely to be blocked by the anti-trust authority. In this case, all firms earn a. This game has four pure-strategy Nash equilibria in which either the merger does not happen (when all firms choose not to join) or a merger between any two companies takes place. More revealing is the fifth Nash equilibrium in mixed strategies. It prescribes that firms choose to join with probability $p = \frac{1}{1+\frac{(c-a)}{2(b-a)}}$. Thus, the incentive to join is inversely related to the ratio of the increase in the payoff of the outsider to the increase in the payoff of the merging parties. The lower is the ratio the higher is the incentive to join. As the results of this paper suggest, risk aversion and/or insufficient downward adjustment of bids by the joint-bidding entity can decrease this ratio substantially.

A natural extension of this work would be to endogenize the decision to bid jointly. This can be accomplished by incorporating a stage where bidders decide whether to participate in a joint-bidding venture or not. A number of interesting questions can be evaluated in this setup. What kind of predictive power does the mixed strategy Nash equilibrium of the above simple game have? Does endogenizing the decision to join have an independent impact on bidders' behavior? Are firms more or less likely to join when the payoff from doing so relative to staying out changes? These questions are left for future work.

To conclude we should mention that other types of value asymmetries in auctions have been recently studied in the literature. Kim and Che (2004) study partitioning of bidders into "knowledge groups" with perfect information about values within a group. Two types of asymmetries arise in this setting: ex-ante asymmetry in group sizes and ex interim asymmetry due to particular realizations of values within groups. Kim and Che (2004) focus on the latter. The former type of asymmetry is very similar to the setting of this paper where a joint-bidding entity can be viewed as a type of knowledge group. There is an important distinction. In the model of Kim and Che (2004) the bidder with the highest value within a group does not bid below the (known) second highest value. In our paper, there is no such constraint as the competition between the members of the joint bidding entity ceases completely. As a result, the anti-competitive effect from the formation of such groups is expected to be larger in the setting of this paper. Another type of ex interim asymmetry is studied in Fang and

Morris (2006). They consider the setup where two bidders in addition to their own private values receive a noisy signal about the rival's value. Such multi-dimensional value structure creates asymmetry in the bidders' perception of the strength of their rivals. Fang and Morris (2006) show that the asymmetry can produce efficient and inefficient symmetric equilibria. The asymmetry is private knowledge, in contrast to the setting of this paper.

6. Acknowledgements

This research has been partially supported by NSF grant number SES-0136928. Any opinions, findings, and conclusions or recommendations are those of the author and do not necessarily reflect the views of the National Science Foundation. I greatly appreciate advice and financial support from John Kagel and Dan Levin. I would like to also thank seminar participants at the Ohio State University and Midwest Economic Theory conference for helpful discussions and useful comments.

References

Andreoni, J., Che, Y., Kim, J., 2007. Asymmetric information about rivals' types in standard auctions: An experiment. Games and Economic Behavior 59, 240–259.

Bornstein, G., Yaniv, I., 1998. Individual and group behavior in the ultimatum game: Are groups more "rational" players? Experimental Economics 1, 101–108.

Chen, K., Plott, C.R., 1998. Nonlinear behavior in sealed bid first price auctions. Games and Economic Behavior 25, 34–78.

Cooper, D.J., Kagel, J.H., 2005. Are two heads better than one? team versus individual play in signaling games. The American Economic Review 95, 477–509.

Cox, J.C., Hayne, S.C., 2006. Barking up the right tree: Are small groups rational agents? Experimental Economics 9, 209–222.

Cox, J.C., Roberson, B., Smith, V.L., 1982. Theory and behavior of single object auctions. Research in experimental economics 2, 1–43.

Cox, J.C., Smith, V.L., Walker, J.M., 1988. Theory and individual behavior of first-price auctions. Journal of Risk and Uncertainty 1, 61–99.

Crawford, V.P., Iriberri, N., 2007. Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in Private-Value auctions? Econometrica 75, 1721–1770.

Fang, H., Morris, S., 2006. Multidimensional private value auctions. Journal of Economic Theory 126, 1–30.

Filiz-Ozbay, E., Ozbay, E.Y., 2007. Auctions with anticipated regret: Theory and experiment. The American Economic Review 97, 1407–1418.

Fischbacher, U., 2007. z-Tree: zurich toolbox for ready-made economic experiments. Experimental Economics 10, 171–178.

Gode, D.K., Sunder, S., 1993. Allocative efficiency of markets with Zero-Intelligence traders: Market as a partial substitute for individual rationality. The Journal of Political Economy 101, 119–137.

Goeree, J.K., Holt, C.A., Palfrey, T.R., 2002. Quantal response equilibrium and overbidding in Private-Value auctions. Journal of Economic Theory 104, 247–272.

Guth, W., Ivanova-Stenzel, R., Wolfstetter, E., 2005. Bidding behavior in asymmetric auctions: An experimental study. European Economic Review 49, 1891–1913.

Hadar, J., Russell, W.R., 1969. Rules for ordering uncertain prospects. The American Economic Review 59, 25–34.

Holt, C.A., 1980. Competitive bidding for contracts under alternative auction procedures. The Journal of Political Economy 88, 433–445.

Huck, S., Konrad, K.A., Muller, W., Normann, H., 2007. The merger paradox and why aspiration levels let it fail in the laboratory. The Economic Journal 117, 1073–1095.

Kagel, J.H., 1995. Auctions: A survey of experimental research, in: John H. Kagel and Alvin E. Roth (eds.), The Handbook of Experimental Economics. Princeton University Press, Princeton, NJ.

Kagel, J.H., Levin, D., 1993. Independent private value auctions: Bidder behaviour in first-, second- and Third-Price auctions with varying numbers of bidders. The Economic Journal 103, 868–879.

Kim, J., Che, Y., 2004. Asymmetric information about rivals' types in standard auctions. Games and Economic Behavior 46, 383–397.

Kocher, M.G., Sutter, M., 2005. The decision maker matters: Individual versus group behaviour in experimental Beauty-Contest games. The Economic Journal 115, 200–223.

Lebrun, B., 1999. First price auctions in the asymmetric n bidder case. International Economic Review 40, 125–142.

Li, H., Riley, J.G., 2007. Auction choice. International Journal of Industrial Organization 25, 1269–1298.

Marshall, R.C., Meurer, M.J., Richard, J., Stromquist, W., 1994. Numerical analysis of asymmetric first price auctions. Games and Economic Behavior 7, 193–220.

Maskin, E., Riley, J., 2000. Asymmetric auctions. Review of Economic Studies 67, 413–438.

Pezanis-Christou, P., 2002. On the impact of low-balling: Experimental results in asymmetric auctions. International Journal of Game Theory 31, 69–89.

Sutter, M., 2005. Are four heads better than two? an experimental beauty-contest game with teams of different size. Economics Letters 88, 41–46.

Sutter, M., Kocher, M.G., Straus, S., 2007. Individuals and teams in UMTS-License auctions. SSRN eLibrary .

Waehrer, K., 1999. Asymmetric private values auctions with application to joint bidding and mergers. International Journal of Industrial Organization 17, 437–452.

Weber, R.A., 2003. "Learning" with no feedback in a competitive guessing game. Games and Economic Behavior 44, 134–144.

7. Appendices

7.1. Experimental Instructions

Purpose: This is an experiment in the economics of decision making. The purpose of the experiment is to explore the decisions you make in a specific economic environment.

There are no tricks, deception or manipulation!!!

Payment: Your payment will depend on the decisions you make, decisions of other participants and chance factors. You may earn a significant amount of money which will be paid to you in cash at the end of the experiment. In addition you will be paid $6 for showing-up on time.

Structure of the experiment: The experiment will consist of a number of repetitions (called periods) of a certain economic situation. These periods are completely independent from one another so that your decisions in one period have no implications for any other period.

Basic setup: You will be participating in a sequence of resale auctions. In these auctions participants bid for the right to resell an item. The person who makes the high bid pays whatever she bid and gets the item. The item is then automatically sold to the experimenter at the bidder's resale price.

Resale prices: Each participant is informed about his/her resale price prior to bidding. The resale price is random and can be anything between 0 and 750. Any number between 0 and 750 is equally likely. Your resale price is completely independent from other bidders' resale prices or your resale prices in previous periods. You can think of an unbiased roulette wheel with numbers from 0 to 750 spun for each participant in every period.

Each auction period will have 3 stages.

Stage 1: In Stage 1 *three* bidders will each bid individually for the right to resell an item. The high bidder gets the item and earns a profit (or loss) equal to the difference between the resale price and what he/she bid:

PROFIT = RESALE PRICE - BID

In other words you bid to buy at your BID and sell at your RESALE PRICE. If you are not the high bidder your profit is zero in this auction. In the case of tied bids, the computer randomly determines who gets the item.

Example*: Suppose three bidders have resale prices 2, 5, and 9 and bid 1, 2 and 7 respectively (these numbers are for illustrative purposes only). Then, the person who bid 7 gets the item because her bid is the highest. She earns a profit of 9-7=2, the difference between the resale price and her bid. The other two bidders get zero profit in this auction. Note that if she bid 9 her profit would have been zero. If she bid 11 she would have incurred a loss of 2. And if she bid 4 her profit would have been 5.*

	Resale	Bid	Profit
Bidder 1	2	1	0
Bidder 2	5	2	0
Bidder 3	9	7	9-7=2
	if B3 bid	9	9-9=0
	if B3 bid	11	9-11=-2
	if B3 bid	4	9-4=5

Stages 2 and 3: In Stages 2 and 3 some bidders will bid in 2-person teams. They will learn each other's resale prices and will bid based on the higher of the two resale prices. That is, if they get the item it will be resold at the higher of their resale prices (which will be referred to as the team resale price). Each auction will have one team and one individual bidder. Members of a team will no longer compete against each other in Stages 2 and 3.

Example*: Suppose three bidders have resale prices 2, 3 and 5. Further suppose the first and the last bidders are in the same team. Their team resale price is therefore 5. The team will submit a bid based on their team resale price of 5 while the remaining bidder will submit a bid based on her resale price of 3.*

Communication in teams: There will be two types of team play. In Stage 2 team members will not be able to communicate with one another and each team member will submit a bid based on the team resale price. In Stage 3 team members will be able to communicate using an instant messaging system and must decide together on a common bid.

Types of bidders: Whether you will be part of a team in Stages 2 and 3 is determined by your type. If your type is TEAM MEMBER you will team up with another bidder. If your type is INDIVIDUAL you will always bid by yourself. The types will be assigned randomly. If you are a member of a team, your partner is the same throughout the experiment.

Matching from period to period: In each auction period teams and individual bidders will be randomly rematched by the computer.

Payoff: Profit is calculated separately for each of the three stages. At the end of each auction period one of these three stages will be randomly selected to be paid off on. The stage selected will apply to all of you in that auction period.

Profits in each stage:

Stage 1: In this stage three bidders compete individually for the right to resell an item. The item is awarded to the highest bidder whose profit is:

PROFIT = RESALE PRICE - BID

The other two bidders get zero.

Stage 2 (when team members cannot communicate with each other): In this stage each team member submits a team bid based on the team resale price. Thus, there are two separate team bids. The individual submits a single bid.

A team member's profit depends on how her **OWN** team bid performs against the bid of the individual. If a team member's bid is higher than the individual bid that team member (and only that team member) earns a profit equal to:

PROFIT = (TEAM RESALE PRICE – TEAM MEMBER'S OWN BID)/2

Otherwise the team member earns zero profit (regardless of her partner's bid).

For the individual we compute her profit against a randomly chosen team bid. If the individual's bid is higher than that team bid her profit is:

PROFIT = INDIVIDUAL RESALE PRICE – INDIVIDUAL BID

Note: if you are a team member your Stage 2 bid affects your profit only. Your partner is not informed about your bid.

Stage 3 (when team members are allowed to communicate): In this stage team members should generate a common bid. If the team bid is higher than the bid of the individual the team gets the item and both team members earn a profit equal to:

PROFIT = (TEAM RESALE PRICE - TEAM BID)/2

If the individual's bid is higher she earns the item and gets a profit equal to:

PROFIT = INDIVIDUAL RESALE PRICE – INDIVIDUAL BID

Converting earnings in the experiment into $: Whatever amount you earn during the experiment will be converted into dollars at the rate of $1 per 40 experimental currency units earned.

Next, we will talk about the software you will use during this experiment.

7.2. Additional Data Tables and Figures

Table 10: Pooled regressions

		RE	OLS (cluster)
Stage 1	*const*	-8.9452	-9.3787**
	v	.93524***	.93402***
	v^2	-.0002***	-.0002***
	s	2.4078	-0.20266
	$s*v$	-0.07922	-0.05964
	$s*v^2$	7.70E-05	5.20E-05
Stage 2	*const*	0.95341	0.95341
	v	-0.01625	-0.01625
	v^2	3.90E-05	3.90E-05
	s	-31.017*	-27.415**
	$s*v$	0.14365	0.1193
	$s*v^2$	-.0002*	-.00017*
Stage 3	*const*	2.9339	2.9339
	v	-0.03346	-0.03346
	v^2	8.50E-05	8.50E-05
	s	-27.501	-23.899**
	$s*v$	0.16301	.13867**
	$s*v^2$	-.00028**	-.00025***
R^2		0.81	0.81

* - p<0.1, ** - p<0.05, *** - p<0.01;
v - value; v^2 - value squared; s - dummy for strong

Table 11: Experimental sessions

Session	Number of subjects Strong	Weak	Omitted	Available data per period per stage #strong bids	#weak bids	#groups
1	14	7	1 weak	14	6	6.5
2	16	8	2 strong	14	8	7.5
3	12	6	0	12	6	6
Total	42	21		40	20	20

Data from one weak subject were omitted as the subject chose to leave the experiment and was replaced by a stand-in subject familiar with the setup (the data from the stand-in subject were omitted too). Data from one strong team were omitted due to a malfunction of the messenger software which precluded communications between the team members for a few periods. The malfunction went unnoticed by the other bidders as the members of the team were still able to submit stage 3 bids (one of which, if different, was randomly chosen as the team's bid).

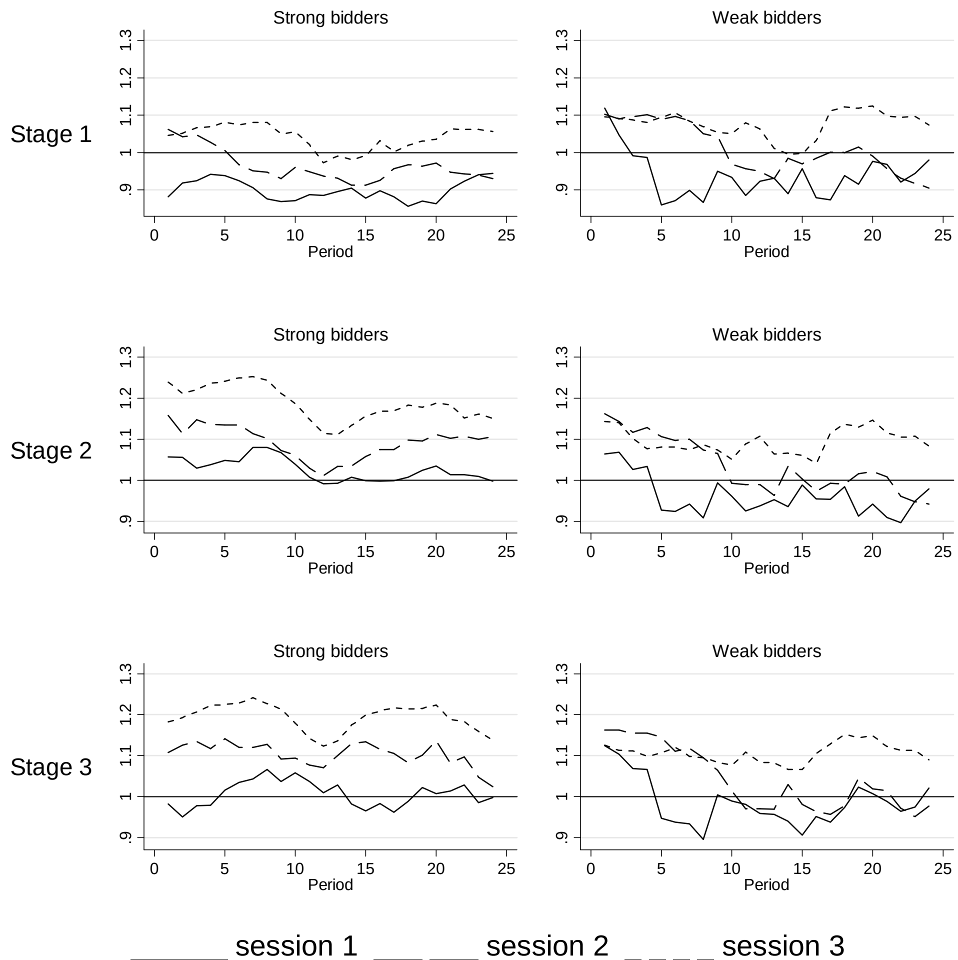

Figure 7: Bidding over time: ratio of observed bids to $\text{RANE}_{0.5}$ bids $(\bar{b}_i/\bar{\beta}_i)$; time series of period averages by session; 3-period moving average smoothing.