Bad and Good News About The Sealed-Bid Mechanism: Some Experimental Results

By ANDREW SCHOTTER*

This paper surveys a set of experiments conducted by Roy Radner and myself (1989) and Peter Linhart, Radner, and myself (1989a; b), all of which were conducted to investigate the properties of the sealed-bid mechanism, a mechanism used to structure bargaining under situations of incomplete information. What we find is both bad news and good news for the mechanism. The bad news is that the behavior of subjects seems quite sensitive to the parameters of the prior distribution of types used and the number of rounds over which the experiment is run. As such prior distributions become more skewed (in a sense to be defined below), and the number of rounds in the experiment are increased to 75, bidder behavior becomes less linear. Such movements have a regular pattern, however, and this phenomenon is remarkably consistent.

The good news is that such behavioral shifts do not seem to interfere with the efficiency of the mechanism. This is of course important for the acceptance of the mechanism in the real world, since if efficiency were to change dramatically as we change the parameters of the environment, adoption would have to proceed on a case-by-case basis. The fact that our results imply robustness of efficiencies indicates that, at least on a practical level, the sealed-bid mechanism may be a viable way to structure bargaining in many large-scale economic organizations.

*New York University, New York, NY 10003. Much of this work has been carried out under NSF grant no. SES-871277 given to New York University. The support of the C. V. Starr Center for Applied Economics is also acknowledged. Finally, I thank Vicky Myroni for writing the programs upon which these experiments were run and Ken Rogoza for valuable research assistance.

I. The Sealed-Bid Mechanism

To quickly summarize the workings of the mechanism, assume that a potential buyer B and a potential seller S are bargaining over the terms of a possible trade of a single object. If the object is traded, the value to B is V and the cost to S is C. (The seller incurs no cost if there is no trade.) The *sealed-bid mechanism* works as follows: B and S simultaneously choose bids, v and c, respectively. If $v \geq c$, then the trade takes place, and B pays S the price $P = (v+c)/2$, that is, the average of the two bids. If $v < c$, then no trade takes place and B pays S nothing.

Suppose that, at the time of bidding, B knows V but not C, and S knows C but not V. The situation is modeled by supposing that V and C are random variables with a joint probability distribution called the *prior*, that is known to both parties. For all experiments surveyed here, it was assumed that the values and costs of the buyers and sellers are drawn independently from the closed interval $[0,100]$ using the following distributions:

$$P(V) = 1 - ([(100-V)/100])^{r_1};$$

$$P(C) = (C/100)^{r_2}.$$

By varying r_1 and r_2 from 0 to 1, we can move these distributions from the perfect certainty case ($r_1 = 0$) to the case of a uniform distribution ($r_i = 1$). When $r_1 = r_2$, we call the mechanism "symmetric," while when $r_1 \neq r_2$, the mechanism is "asymmetric." Before the bidding takes place, B observes V but not C, and S observes C but not V. Buyer B's strategy is a function β that determines his bid v for each value of V, and S's strategy is a function ψ that determines his bid c for each value of C. Second-best optimal linear strategies were derived by

Kaylan Chatterjee and William Samuelson (1983) for environments in which the priors were uniform, while Roger Myerson and Mark Satterthwaite (1983) investigated the welfare properties of this mechanism in more general environments.

There are other equilibria to this mechanism as Wolfgang Leininger et al. (1989) and Satterthwaite and Steven Williams (1989) have shown. In fact, there is an infinite number of other such equilibria, some of which involve continuous but nonlinear bid functions, and some of which involve discontinuous-step functions characterized by flat segments to their bidding functions (segments of slope zero).

The Experiments Performed. While space does not permit even a cursory description of the experiments performed, suffice it to say that subjects engaged in an exercise that faithfully captured the characteristics of the sealed-bid mechanism and motivated subjects to perform with salient payoffs. Except for Radner's and my paper, all experiments were performed on a networked set of PCs. In that paper, in all experiments except one, subjects had uniform priors induced upon them ($r_1 = r_2 = 1$) and repeated the experiment 15 times with the same partner. In the exceptional experiment, a 40-round horizon was used and nonuniform priors induced in which buyers had distribution functions skewed to high values and sellers had distribution functions skewed to low costs ($r_1 = 0.4, r_2 = 0.4$).

II. Radner-Schotter

What Radner and I found was that there was indeed at least qualitative support for the linear equilibrium bid strategies. (Similar linear bid functions were also found by James Cox et al., 1986, for sealed-bid auctions). The only exception to this result was our 40-round experiment in which nonuniform value and cost distributions were used ($r_1 = 0.4, r_2 = 0.4$). In these experiments, a modified form of step-function equilibrium was observed in which the bid functions of the buyer and sellers consisted of a two-piece piecewise linear function with a zero-sloped flat segment followed by a nonzero slope linear segment for the sellers and a nonzero sloped segment followed by a flat segment for the buyers.

The question left unanswered by Radner's and my paper was what caused this behavior shift. Since step-function-like behavior was *only* seen in an experiment where a nonuniform distribution was used *in conjunction with* a 40-period horizon, it was unclear whether it was the horizon or the change in the distributions that accounted for the change.

These findings lead Linhart, Radner, and I (1989a) (hereafter, L-R-S) to perform some further experiments.

III. L-R-S (1989a)

In an effort to answer this question, L-R-S (1989a, b) performed two separate series of experiments that varied the prior distribution coefficients and the number of rounds of the experiments. The results of the experiments are summarized as follows. First, L-R-S were successful in replicating the step-function-like behavior exhibited in Radner's and my paper. This behavior was observed quite frequently and in very dramatic fashion. It is a robust feature of the mechanism in practice. Further, it appears that a 75-round horizon is a necessary but not sufficient condition for this behavior to emerge. To insure quasi-step-function behavior, one must move the r_1 and r_2 coefficients away from 1 and increase the experiments horizon to 75 rounds (or at least past 30 rounds).

Second, the bid and ask functions of subjects seemed to fall naturally into three distinctive types which, for lack of better terms, we have called linear, broken, and bent. While these terms are defined more precisely in L-R-S (1989a), a subject's bid (ask) function is called "linear" if it has no significant changes in slope over its domain. A subject's bid (ask) function is called "broken" if it exhibits one significant change in slope at some point, and if it has a slope of zero over some substantial portion of its domain. A subject's bid (ask) function is called "bent" if it exhibits one significant change in slope at some point but has no segment whose slope is zero.

The frequency with which these types appear is clearly related to the horizon and coefficients of the experiment being performed. For example, in the 15-round experiments run, only 2 out of 31 buyers were categorized as anything but linear. (Both of these occurrences appeared in the experiments where both the buyer and the seller had coefficients of .4 in their distribution functions.) Three out of 31 sellers were characterized that way. Hence, when the horizon of the experiment is 15 rounds (and the coefficients of the distribution function are symmetric), almost all buyer and seller behavior can be characterized as linear. In contrast, when the experiments last for 75 rounds (while the symmetry of the distributions functions is maintained), the distribution of types becomes much more diverse with 12 buyers being categorized as linear, 8 as bent, and 5 as broken. The distribution for the sellers is similarly dispersed with 9 being categorized as linear, 7 as bent, and 9 as broken. Going to 75 rounds seems necessary for the diversification of behavioral types. However, it is not sufficient.

In terms of efficiencies, the sealed-bid mechanism seems to be remarkably robust in its ability to generate efficiency levels that do not vary as either the parameters of the experiment change, or as the subjects alter their bidding behavior. In general, across all subject-pairs efficiencies, when measured as the fraction of the total amount of first-best gains from trade captured, averaged in the 80s. Efficiencies are also consistent over the horizon of the 75-round experiments exhibiting no trend toward either improvement or deterioration when we compare the first, middle, or last 25 rounds. In addition, there is no consistent pattern of efficiency change either as a function of the behavioral types of our pairs or as a function of the experiments parameters. A two-way ANOVA test performed in L-R-S (1989a) substantiates this point. In addition, a series of Mann-Whitney U-Tests uncovered no effect of time, or learning on efficiency in the sense that holding the bidding behavior of pairs constant, the efficiencies observed by subjects using these behaviors did not vary significantly between the first, second, and third 25 rounds of our 75-round experiment. Similarly, using the same techniques, we saw no improvement in efficiency over time when we hold the prior distribution coefficients constant. These results substantiate the claim that the mechanism appears to be robust not only to the parameters of the environment but also to the manner in which people behave under it given these parameters.

IV. Linhart, Radner and Schotter (L-R-S 1989b)

Experimental Designs. While L-R-S (1989a) did yield some of the answers to the questions posed above, there are still some questions left unanswered. First, L-R-S (1989a) leaves unanswered the question of why the behavior change observed in Radner's and my paper appears in the first place? Further, are subjects able to infer from the information given to them during the experiment what type of bid function their partner is using, and is this inference important for them? Finally, given these observed changes, are the subjects acting optimally, that is, if a subject knew that his opponent was employing a broken (or linear) bid function, would the bid function observed be a best response?

To answer these questions, L-R-S (1989b) employed an experimental design in which live subject sellers engaged in 75-round sealed-bid experiments with preprogrammed computerized buyers (in some cases the computerized subjects were sellers) employing one of two fixed strategies. These functions were the following: Linear: $b = .82V$; Broken: $b = V$ if $V \leq 60$ and $b = 60$ if $60 \leq V \leq 100$.

In addition to varying the computerized subjects' bid functions, these experiments were run under two information conditions and using two different sets of priors ($r_1 = r_2 = 1$, and $r_1 = .2$, $r_2 = .4$). In the high-information condition, subjects were told distribution of *bids* that they could expect to observe from their computerized partner's bid function, while in the low-information condition, they were told only the prior distribution of their opponents values (this was the information condition used in the live experiments of Radner and myself, and L-R-S, 1989a). The bid functions used were in some sense representative of those we observed in L-R-S (1989a) while the coefficients used for the prior distributions were the ones which yielded the most dramatic linear and broken behaviors. Clearly, a 75-horizon experiment was also a necessity. Seventeen experiments were performed.

For each computerized bid function, two separate experiments were performed: one each using different prior distributions ($r_1 = .2$, $r_2 = .4$ and $r_1 = 1$, $r_2 = 1$). The order of these presentations was altered as well. Extensive post-experiment surveys were also used. This design does not help in answering the first question posed above, since it only investigates the behavior of subjects in response to their computerized partner's particular linear or broken strategy. It offers no explanation of why *pairs* of subjects change their behavior as the parameters of the environment change.

More precisely, why do subject pairs choose complementary broken strategies when the coefficients in their prior distributions move sufficiently toward zero (given a long enough horizon)? One conjecture is that when it is common knowledge that the buyer's values are sufficiently skewed toward high realizations, and the seller's costs sufficiently skewed toward low realizations (as in our experiments where $r_1 = 0.2$ and $r_2 = 0.4$), subjects treat these situations as the certainty case:[1] $r_1 = 0$, $r_2 = 0$. More precisely, they treat highly skewed cumulative distributions as degenerate distributions with a point mass at 0 for sellers and a point mass at 100 for buyers. In the certainty case, when one chooses a bid, there is no fear that there are no potential gains from trade. In fact, each subject knows that there exist the maximal gains from trade available. The only question becomes one of distributing these gains from trade, so that, over time, we would expect that in a certainty experiment some commonly agreed to price would emerge. The repeated submission of this agreed to bid and ask would take the appearance of a broken-bid function, but, to properly see the analogy, we will have to observe bids as a function of time and not of realized values and costs, since they have a degenerate distribution. Once implicitly agreed to, such a price furnishes a Nash equilibrium.

To investigate this conjecture L-R-S (1989b) ran an additional certainty experiment ($r_1 = r_2 = 0$) which we hoped to use as a basis of comparison for our previously skewed and uniform distribution experiments (1989a). In this experiment, it was common knowledge that buyers would have a value of 100 in every round and sellers would have a cost of 0. This was done with two-sided live subjects since we are interested in the effect of certainty on joint subject behavior.

Results. The L-R-S (1989b) experiments in which live subjects play against computerized partners indicate quite clearly that the bidding behavior of live subjects is not influenced to any significant degree by the bidding strategy of their opponents, but rather only by the parameters of the experiment when their partners are programmed to use specific linear or broken bid functions. By this, we mean that if one compares the results of experiments with identical priors but different bidding functions for the computer, the observed bidding functions for subjects do not appear to differ significantly. However, a comparison experiments in which the bidding strategy of the computer is held constant, but the prior distributions change, does indicate a significant amount of change. While this appears strange at first glance, it is not surprising once one investigates the best-response functions of live subjects facing these computerized subjects.[2] For example, Figure 1 presents the best-response

[1]For example, when $r_1 = .2$ (.4) there is only a 24 percent (43 percent) chance that a buyer will receive a value less than 75 or a 76 percent (57 percent) chance that he will receive a value greater than 75. Likewise, for $r_2 = .2$ (.4) the seller has only a 24 percent (43 percent) chance of getting a cost greater than 25.

[2]Linhart furnished the calculations upon which this statement is made.

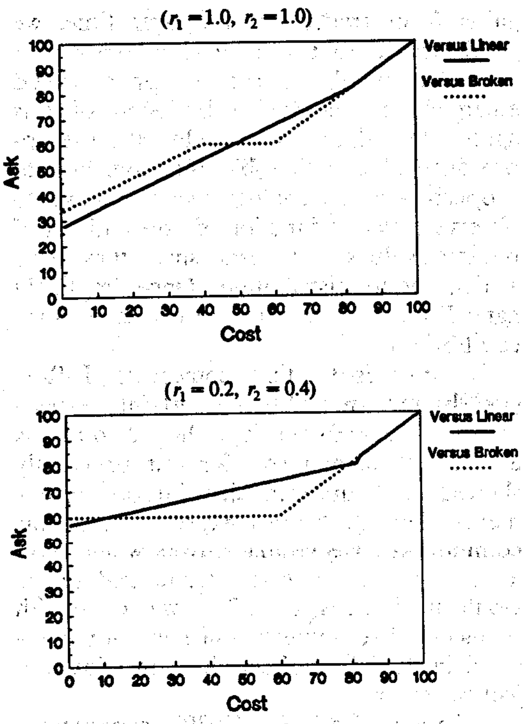

Note: For the best response to the broken strategy, it is assumed that ask = cost for costs greater than 60. Actually, any cost-covering ask is a best response, but asking cost weakly dominates.

FIGURE 1. SELLER BEST RESPONSES VS. LINEAR AND BROKEN BUYER BID FUNCTIONS

function of our subjects to the strategy of their computerized partners as well as the parameters of the experiment's environment. As we see, bidding behavior is far more responsive to changes in the parameters of prior distributions than changes in the bidding behavior of their opponent. At least, on a qualitative level, it is shown in my paper (1990) that the mean actual behavior of subjects did in fact conform to these best-response functions.

The final conjecture of L-R-S (1989b) concerned the cause for the shift in bidding behavior as the parameters of the experiment changed. It was conjectured that in those experiments where the prior distributions were sufficiently skewed, subjects would behave as they would in the certainty case.

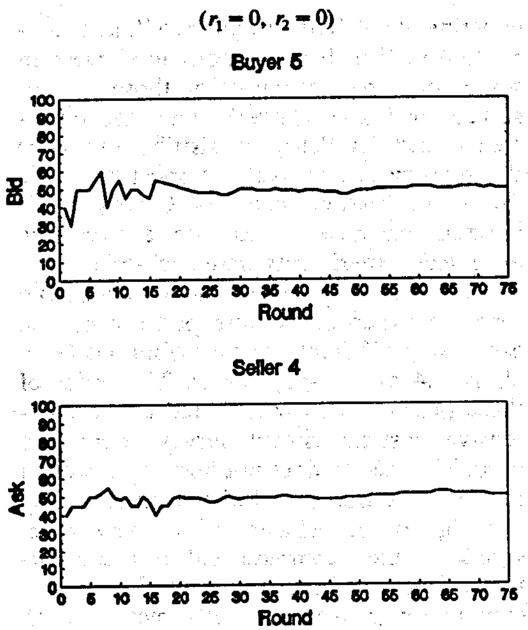

FIGURE 2. BID-ASK FUNCTIONS OVER TIME: THE CERTAINTY CASE

In the certainty case, we expect to see an agreed-to splitting of the gains from trade. Hence, if we plot the bids made by subjects against time, we would expect to see eventually both buyer and seller converging on the same agreed upon price. Figure 2, which presents one example of bid and ask time patterns for one buyer and one seller, substantiates this belief. (Space limitations require that I present only one example. This example is representative of the vast majority of cases, however.)

While in some cases there is a trend exhibited, in general, a constant bid function appears. Equivalent (certainty-like) behavior in the skewed experiments ($r_1 = 0.2$, $r_2 = 0.4$) would generate bid (or ask) functions that, after a certain amount of experimentation, were flat with respect to time, but which exhibited upward spikes (downward spikes) whenever a value (or cost) is realized that is above (below) the agreed upon price. Such behavior would be consistent with our two-piece piecewise linear function projected onto in value bid (cost-ask) space. Figure 3 pre-

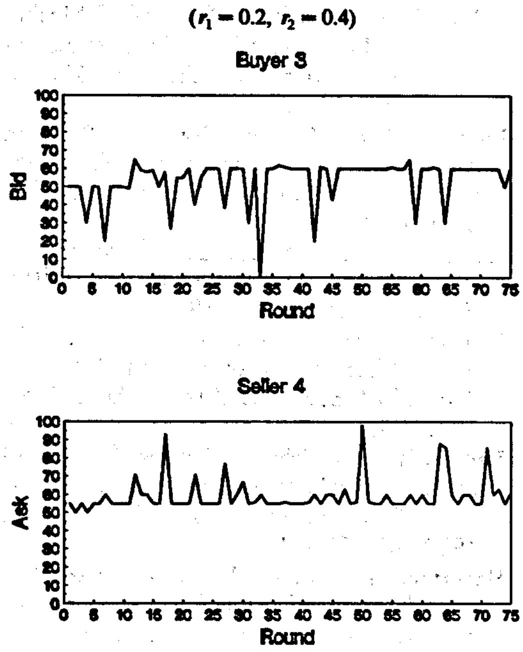

FIGURE 3. BID-ASK FUNCTIONS OVER TIME: THE SKEWED DISTRIBUTION

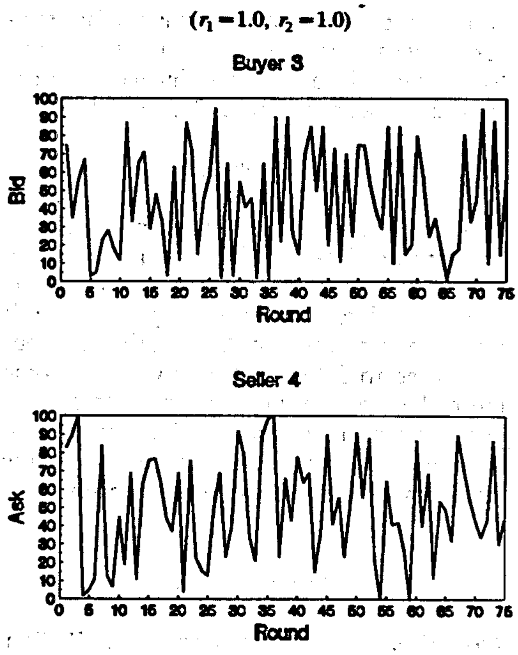

FIGURE 5. BID-ASK FUNCTIONS OVER TIME: THE UNIFORM DISTRIBUTION

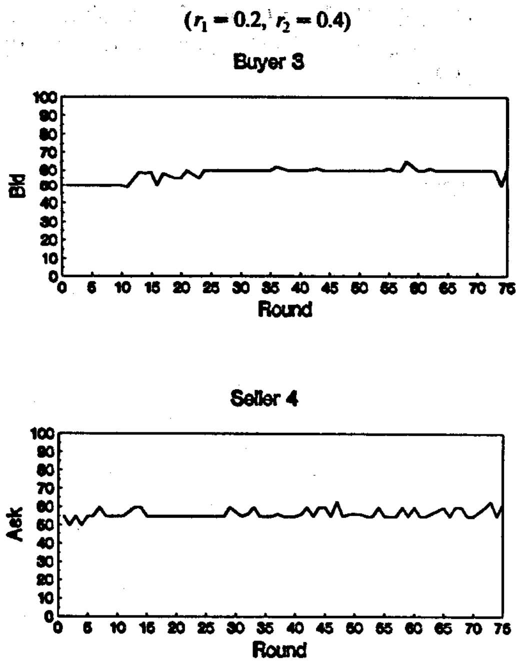

FIGURE 4. BID-ASK FUNCTIONS OVER TIME: THE SKEWED DISTRIBUTION MAJOR SPIKES REMOVED

sents another representative sample of the bid and ask time patterns in the $r_1 = 0.2$, $r_2 = 0.4$ live experiments of L-R-S (1989a). Note that these functions are basically constant with respect to time except for spikes that occur when a value or cost is realized above or below the implicitly agreed to price. These spikes are, of course, necessary in order to avoid losses on the parts of subjects. If we were to eliminate these spikes, we would create Figure 4 which has much of the appearance of the graphs in Figure 2. The comparison of Figures 2 and 4 help substantiate the conjecture that sufficiently skewed distributions may be treated like degenerate certain distributions by subjects.

These diagrams are in sharp contrast to the type of behavior observed in those experiments where all prior distributions are uniform ($r_1 = r_2 = 1$). In these experiments, if subjects bid some linear function of their value or cost, we would expect to see a time graph of bids which was simply a linear transformation of the realizations of these random variables. In Figure 5 we present another representative example of a bid-and-ask pattern from the uniform experiments. Note that this behavior is indeed different from those seen in Figures 2, 3, and 4.

V. Conclusions

What have we learned about the sealed-bid mechanism? First, we know that depending on the environment in which it is used, the behavior of experimental subjects using it may vary dramatically. As the prior distributions used to define costs and values become more skewed, the bidding functions of subjects become more broken. Such changes in bidding behavior seems not to affect the efficiency of the sealed-bid mechanism.

REFERENCES

Chatterjee, Kaylan and Samuelson, William, "Bargaining Under Incomplete Information," *Operations Research*, 1983, *31*, 835–51.

Cox, James, Smith, Vernon and Walker, James, "Theory and Individual Behavior of First Price Auctions," unpublished manuscript, University of Arizona, May 1986.

Leininger, Wolfgang, Linhart, Peter and Radner, Roy, "Equilibria in the Sealed-Bid Mechanism for Bargaining With Incomplete Information," *Journal of Economic Theory*, June 1989, *48*, 63–106.

Linhart, Peter, Radner, Roy and Schotter, Andrew, (1989a) "Behavior and Efficiency in the Sealed-Bid Mechanism: Part 1," mimeo., 1989.

______, ______, and ______, "Behavior and Efficiency in the Sealed-Bid Mechanism: Part 2," in progress, 1989.

Myerson, Roger and Satterthwaite, Mark, "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, April 1983, *29*, 265–81.

Radner, Roy and Schotter, Andrew, "The Sealed-Bid Mechanism: An Experimental Study," *Journal of Economic Theory*, June 1989, *48*, 179–221.

Satterthwaite, Mark and Williams, Steven, "Bilateral Trade With the Sealed-Bid *k*-Double Auction: Existence and Efficiency," *Journal of Economic Theory*, June 1989, *48*, 107–33.

Schotter, Andrew, "Bad and Good News About the Sealed-Bid Mechanism: Some Experimental Results," working paper no. 90-0, C. V. Starr Center for Applied Economics, New York University, January 1990.