ECONOMIC RESEARCH REPORTS

BAD AND GOOD NEWS ABOUT THE SEALED-BID MECHANISM: SOME EXPERIMENTAL RESULTS

BY

Andrew Schotter

R.R. # 90-04 January, 1990

C. V. STARR CENTER FOR APPLIED ECONOMICS

NEW YORK UNIVERSITY
FACULTY OF ARTS AND SCIENCE
DEPARTMENT OF ECONOMICS
WASHINGTON SQUARE
NEW YORK, N.Y. 10003

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

by

Andrew Schotter
New York University

November 1989

This paper was written for presentation at the American Economic Association Meetings in Atlanta Georgia, in December 1989.
This paper surveys the joint research conducted by the author with Roy Radner and Peter Linhart. Much of this work has been carried out under grant number SES 871277 given to New York University by the National Science Foundation. The support of the C.V. Starr Center for Applied Economics is also acknowledged. The author would finally like to thank Vicky Myroni for writing the programs upon which these experiments were run and Ken Rogoza for his valuable research assistance.

Abstract

In this paper we survey a set of experiments performed by Radner and Schotter and Linhart, Radner, and Schotter which investigate one particular way to structure bargaining using what has been called the sealed–bid mechanism.

The authors were interested in investigating the performance characteristics of the sealed–bid mechanism in an effort to see whether it might be a potential mechanism to use in large scale economic organizations. What they 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 were 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.

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

Andrew Schotter

Section 1: Introduction

All price formation involves bargaining. The only difference is in how the bargaining is structured. For example, while most laymen consider bargaining to be only that situation in which two people or more sit face to face and negotiate a price for a given commodity or service, economists recognize that even a posted price retail market in which sellers post a price that is not negotiable and buyers either buy or do not buy is, in essence, an abbreviated form of bargaining in which the rules specify only one exchange of messages (an ask by the seller and a yes or no by the buyer). Given this description of price formation, a natural question arises as how to best structure bargaining so as to facilitate surplus maximizing exchange. In this paper I will survey a set of experiments performed by Radner and Schotter (RS 1989), and Linhart, Radner, and Schotter (LRS (1989a), LRS (1989b in progress)) which investigate one particular way to structure bargaining using what has been called the sealed–bid mechanism.[1]

This mechanism is of interest for two different reasons. First, it is a mechanism that has generated a fair amount of interest recently by economists since under certain circumstances one of its equilibria (the linear equilibrium) is second–best welfare optimal. In fact this "linear equilibrium" generates a higher level of ex–ante gains from trade than

---

1

This limiting of scope will not permit me to include the early and interesting experimental work of William Samuelson (1985) on a variant of the sealed–bid mechanism in which one side makes a take–it–or–leave–it offer to the other. This omission is unfortunate but necessary since the results of Samuelson do not comment on the main theme of the results discussed below) the equilibrium of any other mechanism. As a result it is of interest to discover how this mechanism actually performs in the laboratory to see if it is a candidate for implementation in the real world. Second, as Satterthwaite and Williams (1989) and Wilson (1985) point out, if one were to generalize this sealed-bid mechanism to the case where many units of a homogeneous good were being exchanged between a large number of buyers and sellers, one would see a mechanism which has many of the characteristics of what perfectly competitive price formation might look like if that process were ever rigorously defined. Hence the study of sealed-bid bargaining may be a vehicle through which we might heighten our understanding of competitive price formation. The purpose of the experiments I will survey was motivated by the former of these two reasons. That is, the authors were interested in investigating the performance characteristics of the sealed-bid mechanism in an effort to see whether it might be a potential mechanism to use in large scale economic organizations.

What they 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 were 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.

In the remainder of this paper I will proceed as follows. First in Section 2 I will quickly summarize the sealed–bid mechanism and its equilibria. This will be followed in Section 3 by a review of the Radner–Schotter (1989) experimental results on the mechanism, a review which will emphasize what appears to be a surprising shift in the behavior of subjects as the parameters of that experiment are varied. Section 4 will explain the design and results of two further studies (LRS (1989a) and LRS (1989b)) whose aim was to try to understand what was observed in RS (1989a). Section 5 offers some concluding remarks. Because this paper is a survey of three rather complicated sets of experiments, and due to space limitations, I will not get involved too deeply in the details or statistical procedures used to test our findings and will rely heavily on simple graphs and tables to make my point. While this may tend to make the presentation anecdotal at times, interested readers are referred to LRS (1989a) and LRS (1989b) and RS (1989) for the details about our experimental methods and statistical techniques.

Section 2: The Sealed–Bid Mechanism

To quickly summarize the workings of the mechanism, assume that a potential buyer, B, and a potential sellers, 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 $P = (v+c)/2$, i.e. 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 modelled by supposing that V and C are random variables with a joint probability distribution called the prior, which 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 ($r_i = 0$) to the case of a uniform distribution ($r_i = 1$). When $r_1 = r_2$ we will 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. 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.

If P(V) and P(C) are uniform distributions ($r_1 = r_2 = 1$ in our experiments) defined over the closed interval [0,100], then as Chatterjee and Samuelson (1983) have demonstrated, there exists a pair of linear bidding strategies which together form an equilibrium of the game defined by the sealed-bid mechanism. Myerson and Satterthwaite (1983) have demonstrated that this particular equilibrium has a very strong welfare property: It maximizes the ex-ante gains from trade that can be achieved at any Bayesian-Nash equilibrium of any individually rational bargaining mechanism employed in this environment, in which what the seller received equals what the buyer pays.

There are other equilibria to this mechanism as Leininger, Linhart, and Radner (1989) and Satterthwaite and Williams (1989) have shown. In fact there are an infinite number of other such equilibria some of which involve continuous but non-linear bid functions and some of which involve discontinuous step-functions. In a two-step equilibrium, there is some common value (cost) x such that if the buyer receives a realization greater than x he bids x and if he receives less than x he bids 0. Similarly, the seller, <u>using the same x</u>, bids x for all cost realizations below x and bids 100 for cost realizations above it. Hence, step function equilibria are characterized by flat segments to their bidding functions (segments of slope zero).

2.1: 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 RS (1989), all experiments were performed on a networked set of PC's. In all RS (1989) 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 non–uniform priors induced in which buyers had distribution functions skewed to high values and sellers had distribution functions skewed to low costs ($r_1 = 0.4$ $r_2 = 0.4$). Further explanations of the experimental design used by LRS (1989a and 1989b) are offered in section 4.

Section 3: Radner–Schotter (1989)

3.1: The Results

What RS 1989 found was that there was indeed at least qualitative support for the linear equilibrium bid strategies. (Similar linear bid functions were also found by Cox, Smith and Walker (1986) for sealed bid auctions). The only exception to this result was their 40 round experiment in which non–uniform 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 piece–wise linear function with a zero sloped flat segment followed by a non–zero slope linear segment for the sellers and a non–zero sloped segment followed by a flat segment for the buyers. Figure 1 presents a representative sample of such bidding patterns.

These figures illustrate a type of behavior that was found in this experiment that failed to exist in any of the other RS (1989) experiments. As you can see, the behavior exhibited is a hybrid type of behavior combining segments of linear behavior with that of step–like behavior.

While this pattern was seen in abundance, it does not form an equilibrium to the static game defined by the sealed bid mechanism. The question left unanswered by the RS

Figure 1
Step-Function-Like Bid-Ask Functions, RS (1989)

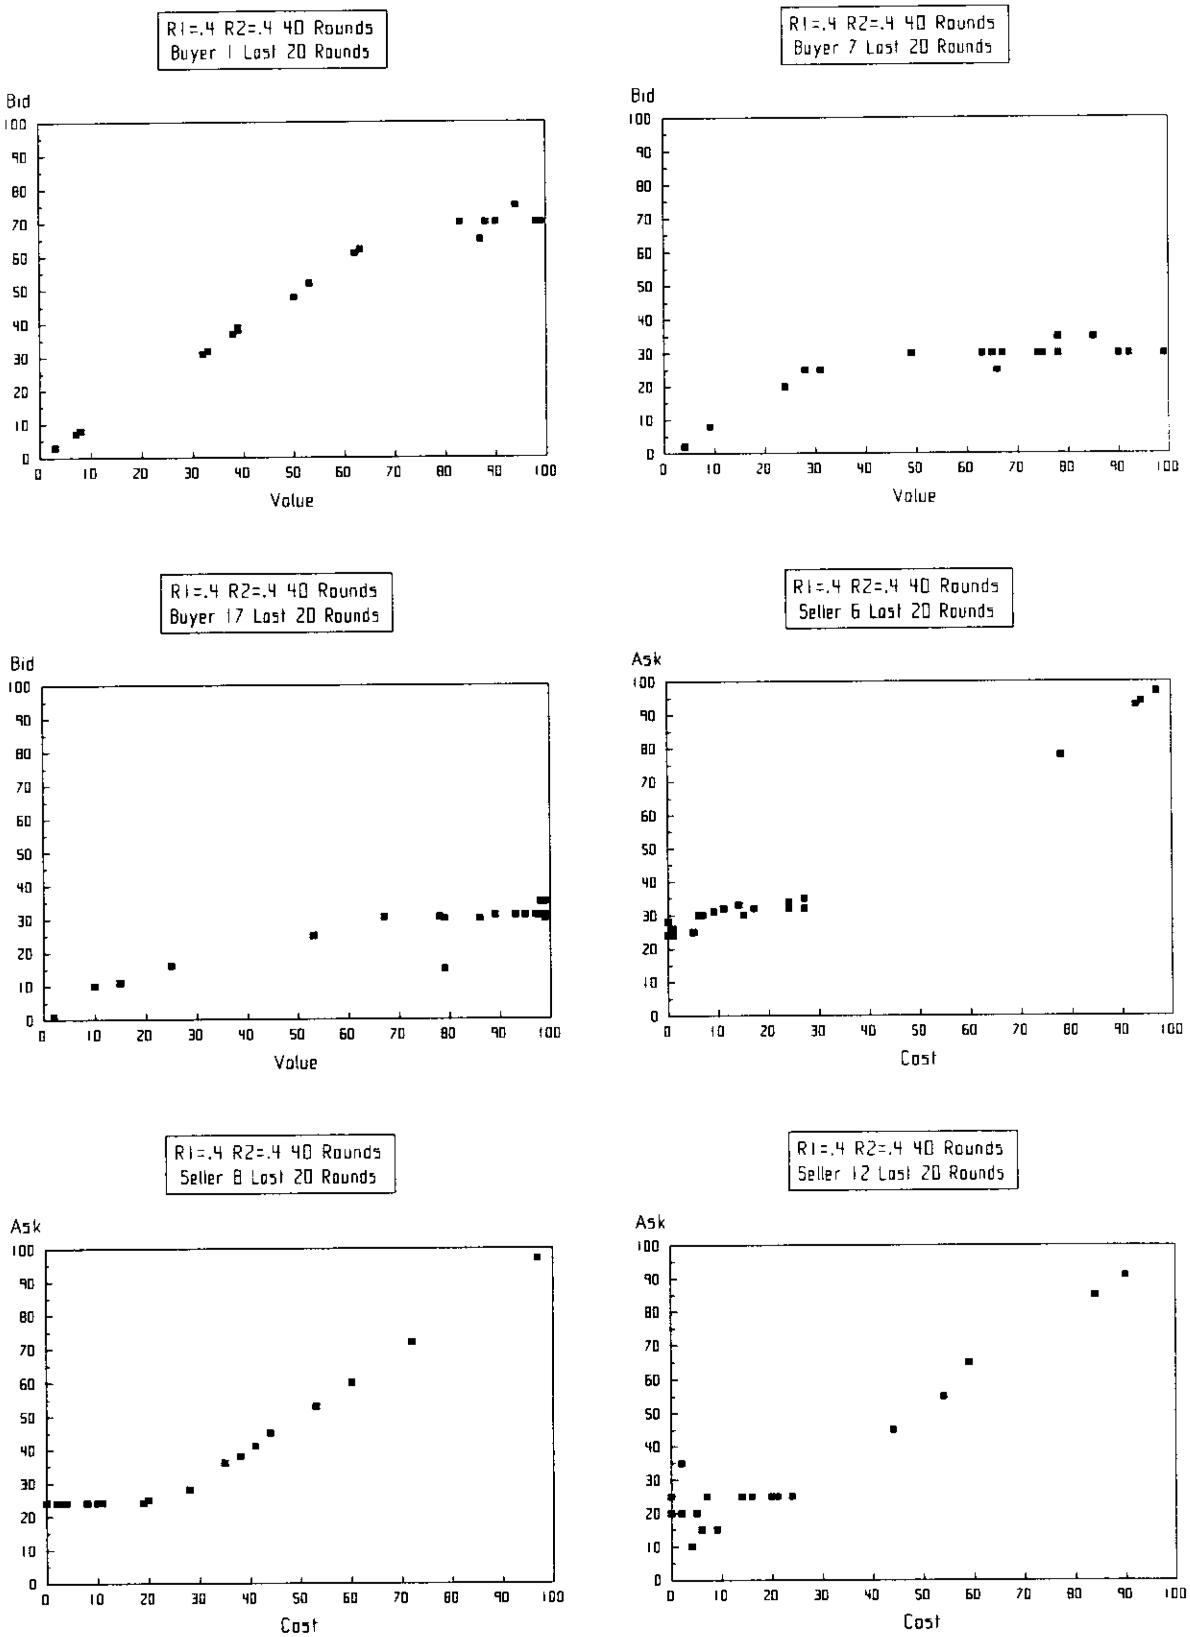

1989 paper was what caused this behavior shift. Since step–function–like behavior was only seen in an experiment where a non–uniform distribution was used in conjunction with a 40 period horizon, it was unclear whether it was the horizon or the change in the distributions that was accountable for the change. For example, was their failure to observe such behavior in their 15 round uniform distribution experiments an artifact of the fact that due to the short horizon subjects did not have a chance to coordinate their strategies as called for by the step–function–like strategies found in the 40 round experiment, or was the behavior change merely a result of changing the distributions in the experiment?

These findings lead Linhart Radner and Schotter (1989a) to ask the following questions:

1) What factors are responsible for the shift in behavior observed in the RS (1989) experiments? In particular, is it the time horizon of the experiment or the coefficients of the distribution functions described above.

2) If it is the coefficients of the distribution function, which coefficient is most important –– the buyers distribution coefficient, the seller's coefficient, neither, or both ?

3) What are the major determinants of the realized efficiencies of the mechanism? Is it the parameters of the experiment or is the bidding behavior of the subjects more important?

4) Do the bidding behaviors of subjects in these experiments fall into recognizable categories and if so what types of behaviors lead to the most efficient results?

5) Are the efficiencies observed in the mechanism sensitive to the recognized bidding patterns of subjects or is the mechanism robust to these changes?

6) How do subjects learn over the course of the experiment? Do they converge on one type of behavior depending on their early experience, and if so what types of experience dictates behavior?

These six questions motivated the experimental design for Linhart, Radner and Schotter(1989a).

Section 4: Linhart, Radner and Schotter (1989a)

4.1: Experimental Designs

In an effort to answer the 6 questions listed above, LRS (1989a) and LRS (1989b) performed two separate series of experiments. In the first (Linhart, Radner and Schotter (1989a) the experimental design is presented in table 1:

Table 1: Experimental Design

r coefficient for Buyers (r_1)	r coefficient for Seller (r_2) 0.2 A	0.2 B	0.4 A	0.4 B	1.0 A	1.0 B
0.2	15 30 75	8 7 6	75	9	75	10
0.4	75	9	15 30 75	12 11 11	75	12
1.0	75	12	75	10	15 30 75	11 6 9

Note: The numbers under the A column in each cell represent the number of rounds to be performed while the numbers under the B column in each cell represent the number of subject pairs used.

This table has a simple explanation. In total, 15 experiments were performed each of which was indexed by a pair of coefficients for the buyer and seller distributions, and a number of rounds. Hence each cell of this matrix locates an experiment performed since the row and column indices locate the coefficients used, while the numbers in column A indicate how many different experiment lengths were used for any given configuration of coefficients. Notice that along the diagonal, when the coefficients were equal, we ran experiments of length 15, 30, and 75, while off the diagonal, where the buyers and seller coefficients were different, only experiments of length 75 were run. The numbers in column B in each cell indicate the number of subject pairs used for that experiment. Subjects were never re–used and only participated in one experiment.

Note that this design allows answers to the questions posed above. For example, if one thinks that it is the horizon of the experiment and not the coefficients of the prior probability distributions that influence behavior, one can investigate the experiments run along the diagonal where <u>only</u> the horizon changes and the coefficients of the priors are held constant. On the other hand, if one feels that it is the coefficients that are important, then one can look across the rows or down the columns and compare experiments of equal 75 round lengths which differ only by their experimental prior distribution coefficients. If one hypothesizes that it is the buyer coefficient (or seller) alone that is responsible for the behavior shift, one can compare experiments of equal length and seller coefficients but with varying buyer coefficients.

<u>4.2 :Results</u>

The results of the LRS (1989a) experiments are easily summarized as follows. First, LRS were successful in replicating the step–function–like behavior exhibited in RS (1989). 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 a 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 leat 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 LRS (1989a), a subject's bid (ask) function is called "linear" if it has no significant changes in slope over its domain. A subject bid (ask) function is called "broken" 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). 3 out of 31 sellers were characterized that way. Hence when the horizon of the experiment is 15 round (and the coefficients of the distribution function are symmetric), almost all buyers 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.

These comments are illustrated by table 2 and figures 2a and 2b (they are also supported by a series of ANOVA tests which we will not report here (see LRS 1989a). Table 2 presents the fraction of buyers and sellers using linear bid functions in each of our experiments while figures 2a (2b) presents the "average" bid (ask) function used by subjects in our experiments. By "average" we mean the bid function derived by fitting the best two–piece–wise linear function to the data for each subject after appropriately eliminating outliers (see Linhart, Radner and Schotter (1989a) for details of this procedure) and then averaging the slopes of these functions to the right and left of their break–point.

As can be seen from table 2 when the horizon of the experiment is fixed at 15 rounds, almost all the buyers and sellers seem to bid in a linear manner. Movement away

Table 2

Buyer and Seller Types Disaggregated by Experiment:
Fraction Exhibiting Linear Bid Functions

			SELLERS							
		r_2	1.0 15	1.0 30	1.0 75	0.4 15	0.4 30	0.4 75	0.2 15	0.2 30	0.2 75
	r_1										
	1.0	15	(11/11) (10/11)								
	1.0	30		(6/6) (2/6)							
	1.0	75			(6/8) (4/8)			(5/10) (7/10)			(2/12) (6/12)
	0.4	15				(10/12) (10/12)					
BUYERS	0.4	30					(8/11) (8/11)				
	0.4	75			(3/11) (4/11)			(3/11) (3/11)			(4/9) (2/9)
	0.2	15							(8/8) (8/8)		
	0.2	30								(7/7) (7/7)	
	0.2	75			(4/10) (4/10)			(0/9) (3/9)			(3/6) (2/6)

Note:
In each cell the format is as follows:
(Fraction of buyers using linear bid function)
(Fraction of Sellers using linear ask function)
[Two pairs were excluded because they could not be classified].

Figure 2a
Average Bid Functions LRS (1989a) - Buyers

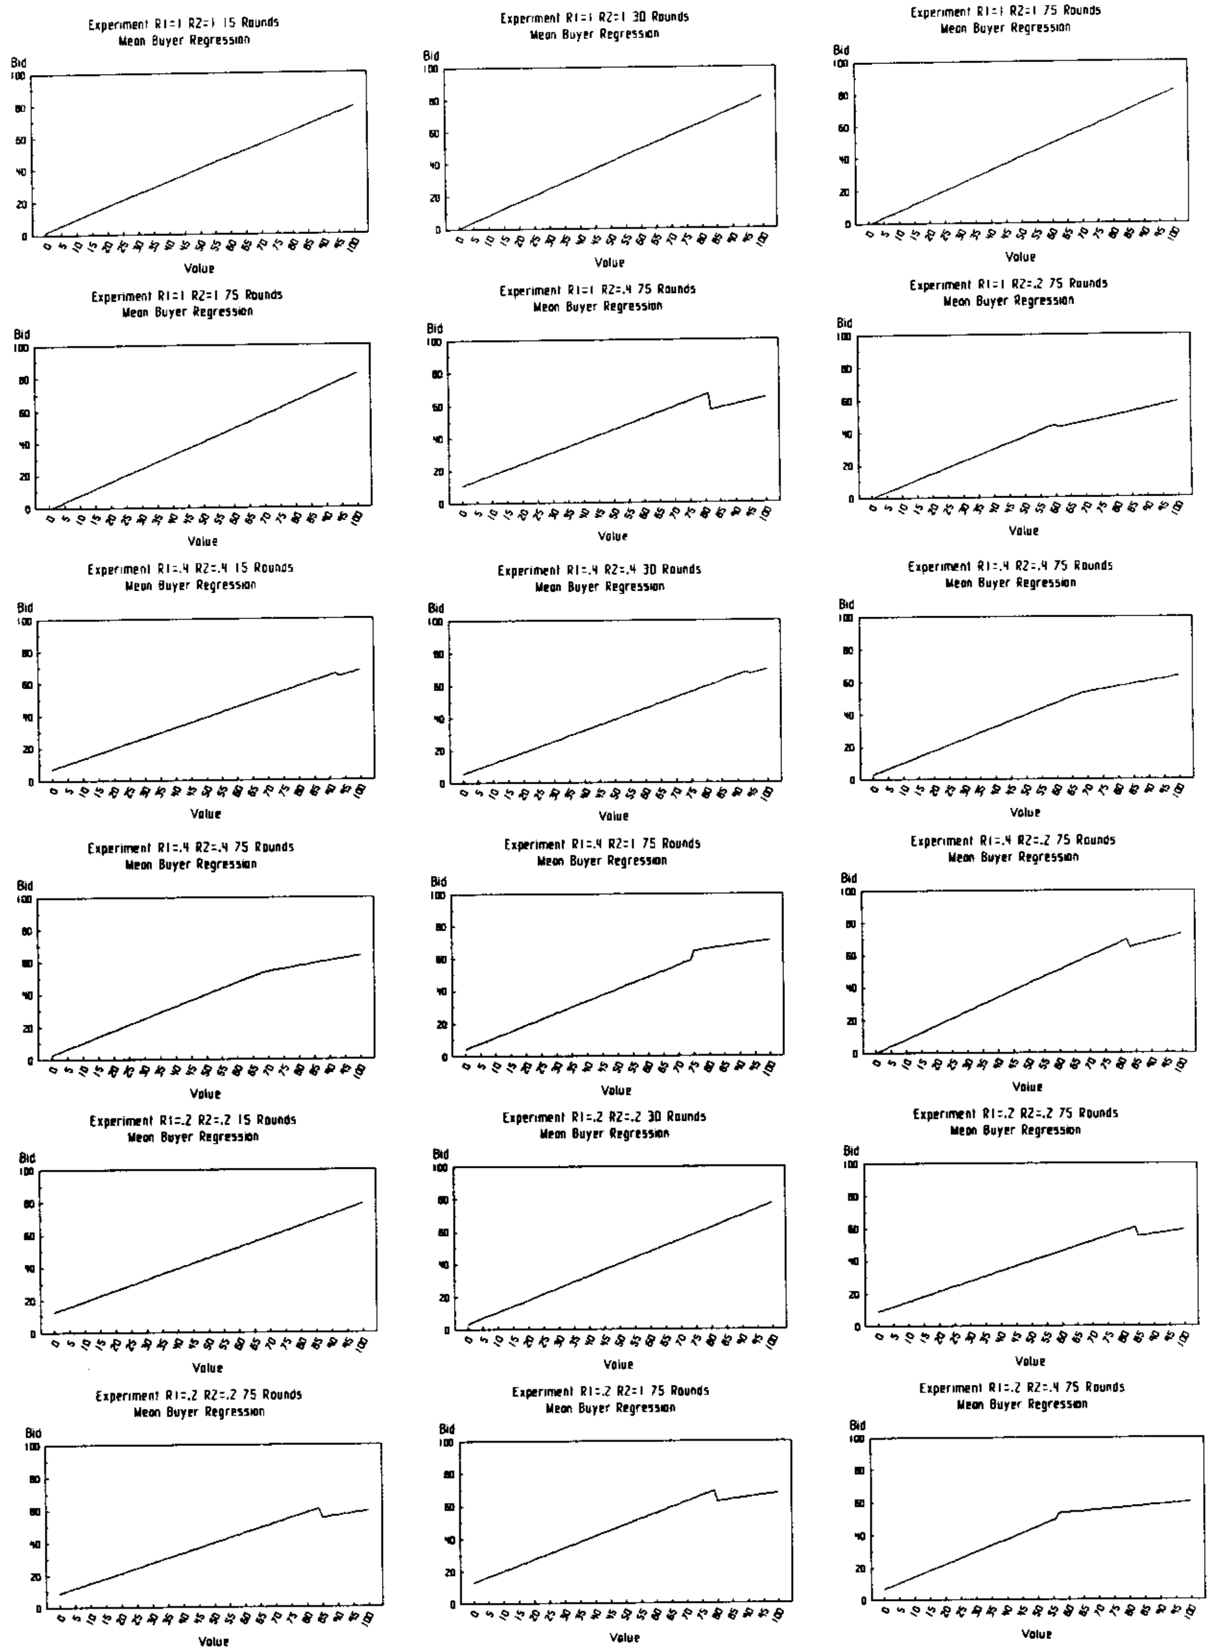

Figure 2b
Average Ask Functions LRS (1989a) - Sellers

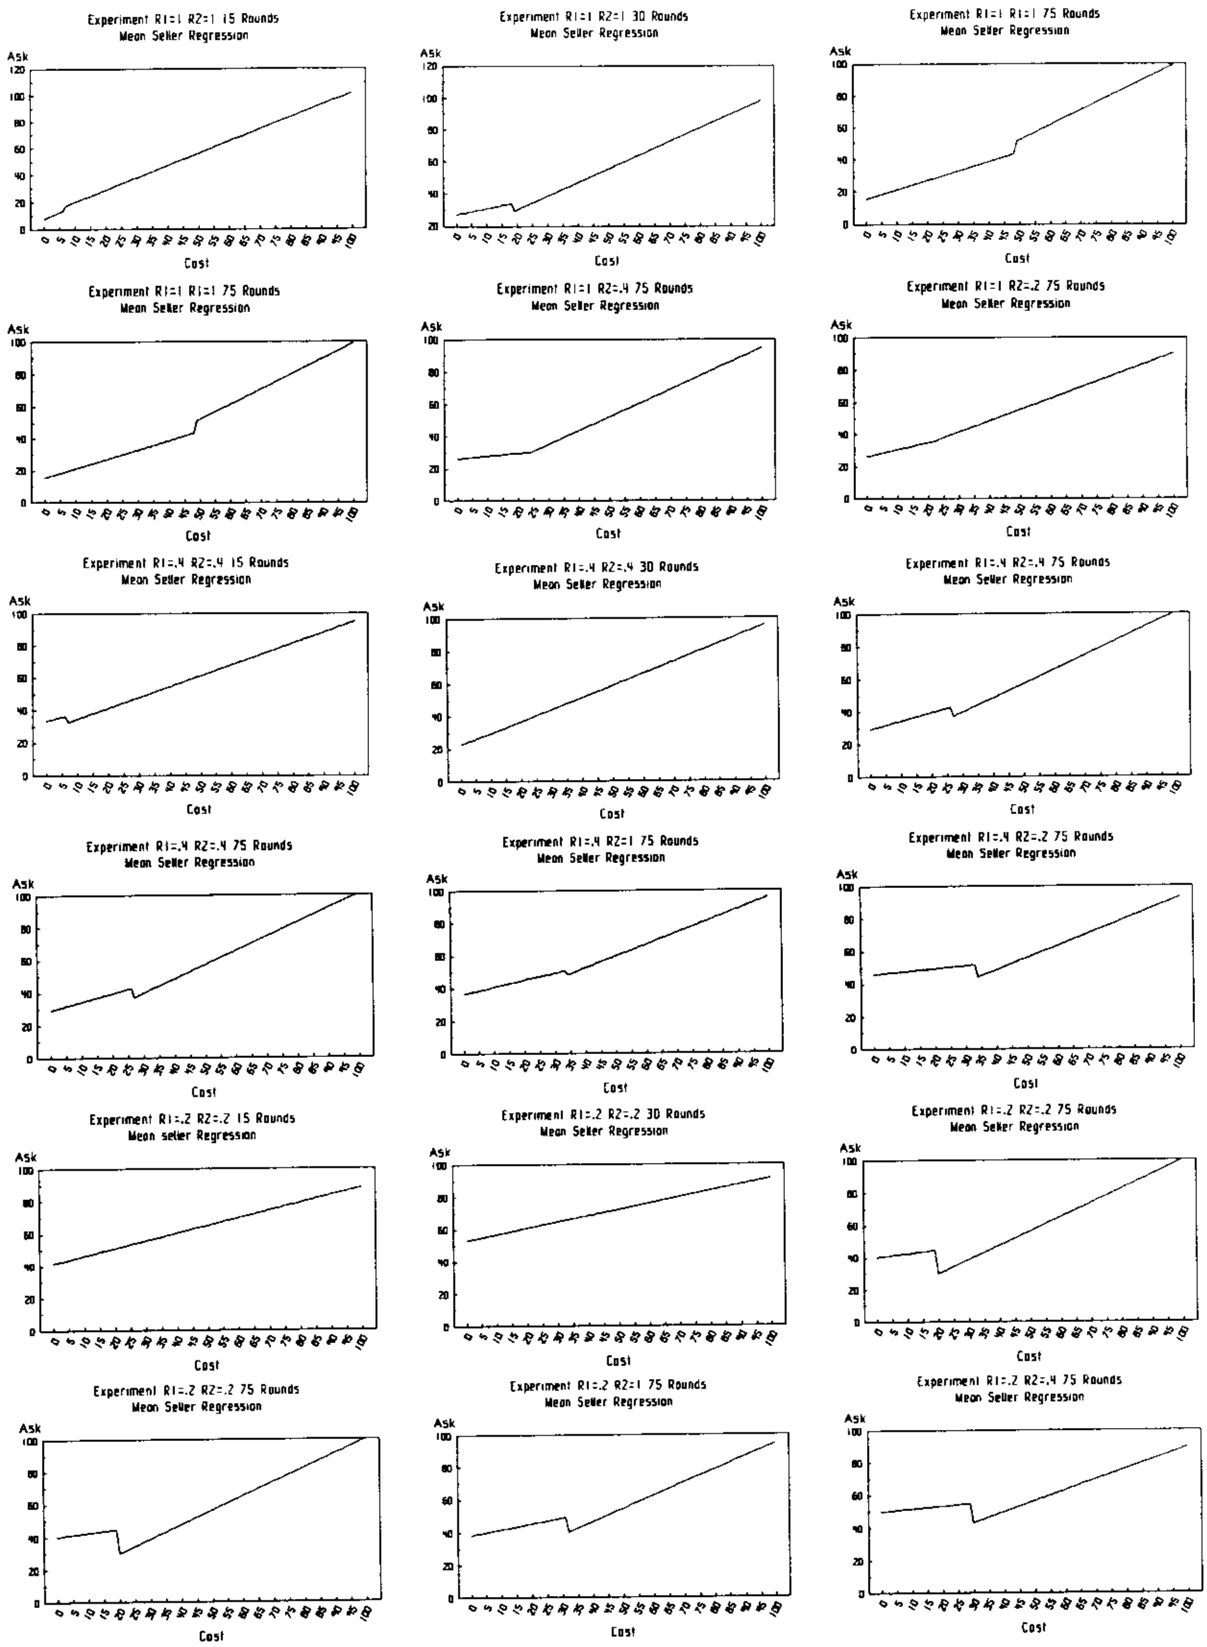

from linearity does not occur in a meaningful way until we reach the 75 round experiments and even then linearity is still a popular bidding pattern among buyers and sellers whose priors are both uniform ($r_1 = r_2 = 1$).

Figures 2a and 2b tell the same story. Note that these average bid (ask) functions are linear (or almost linear) for all experiments with horizons of 30 or less no matter what prior distributions are used, but become bent or broken when the 75 round horizon is reached. The exception to this rule is the experiment where both subjects' priors are uniform, in which case linearity is maintained. It is in this sense that we consider the horizon only to be a necessary condition for step–function–like behavior. When priors are uniform, no length of horizon seems capable of altering the linear nature of subjects' bidding behavior.

In terms of efficiencies the sealed bid mechanism seems to be remarkably robust in its ability to generate efficiency levels which 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 80's. 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. This result is surprising since one would expect that since our bidders exhibited a wide variety of bidding behavior which varied regularly across experiments, we would expect such changes in their behavior to manifest itself in differing efficiencies. An analysis of variance performed on the data rejects this hypothesis. These efficiency results are presented in table 3 where we present the fraction of the first–best gains from trade achieved by pair type over the first, second and final thirds of our 75 round experiments. As we can see, 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 LRS (1989a)

Table 3:

Fraction of the First Best Gains From Trade
Achieved by Pair Bid Type and Experiment
First, Middle, and Last 25 Rounds

	Experiment									
Pair-Type	r1r1	r1r2	r1r4	r4r1	r4r4	r4r2	r2r1	r2r4	r2r2	Mean
Linear-linear	.946 4	.863 5	.805 4	.843 5	.812 2	.874 2	.821 2	.835 3	-	.847 27
	.974 2	.729 1	.778 1	-	.990 1	.853 1	-	.562 1	.912 2	.836 9
	.805 4	.874 4	.912 4	.912 4	.824 2	-	-	-	.947 1	.887 19
Linear-Bent	-	-	.877 1	-	.889 2	-	.721 2	.795 1	.838 1	.827 7
	.804 2	-	-	-	.643 1	-	-	.895 2	-	.809 5
	-	-	-	.924 1	.932 1	.949 1	.876 2	.791 1	-	.894 6
Linear-Broken	.888 3	.739 3	-	.758 2	.770 2	-	.805 3	-	.757 1	.780 14
	-	.764 1	.932 1	.834 1	-	.908 2	.752 2	.875 1	-	.849 8
	-	.756 1	.887 1	-	.747 1	.906 2	.946 1	-	.828 1	.856 7
Broken-Broken	-	.750 1	-	-	-	.908 2	-	.910 1	.913 1	.889 5
	-	-	-	.917 1	.972 1	-	.878 2	.841 1	-	.897 5
	-	-	-	.784 1	.833 1	-	.854 2	-	.825 2	.828 6
Broken-Linear	-	-	-	.856 1	.874 2	-	.932 1	.751 1	.732 3	.732 8
	-	.719 1	-	.899 1	.898 3	-	.912 1	-	.504 1	.777 7
	-	.921 2	.892 1	.761 1	.923 4	-	.894 2	.929 1	.897 1	.906 12
Broken-Bent	-	-	-	.781 1	-	.869 1	.898 1	-	-	.857 3
	-	-	.822 1	.838 3	-	.959 1	-	.815 1	.841 1	.855 7
	-	-	-	.897 3	-	.885 3	-	.876 2	-	.886 8

Table 3 (continued)

	Experiment									
	r1r1	r1r2	r1r4	r4r1	r4r4	r4r2	r2r1	r2r4	r2r2	Mean
Bent-Bent	- 	.923 1	.868 1	.857 2	.712 1	.912 2	- 	.904 1	- 	.872 8
	.822 3	- 	.808 1	- 	.833 2	.887 3	.890 1	.876 1	- 	.863 11
	- 	- 	.854 2	- 	- 	- 	.794 1	- 	- 	.838 3
Bent-Linear	.480 2	.823 2	.811 3	- 	.866 2	.730 2	.910 1	.793 1	- 	.793 13
	.753 2	.835 9	.864 5	.822 3	- 	.965 1	.917 2	.941 1	.946 1	.868 24
	.854 5	.698 4	- 	- 	.830 2	.896 2	.679 1	.877 2	- 	.813 16
Bent-Broken	- 	- 	- 	- 	- 	- 	- 	.903 1	- 	.903 1
	- 	- 	- 	.814 2	.861 3	.902 1	.703 2	.816 1	.837 1	.827 10
	- 	.709 1	.908 1	.737 1	- 	.867 1	.788 1	.836 2	.826 1	.822 8

Note: Efficiencies for two pairs are not reported since they could not be classified. An additional efficiency is not reported for the last 25 rounds as one pair had no opportunites to trade in those rounds.
Note: All cells of this matrix are read as follows;

Efficiency 1^{st} 25 rounds,
observations
Efficiency 2^{nd} 25 rounds
observations
Efficiency 3^{rd} 25 rounds
observations

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 they 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.

4.3: Linhart, Radner and Schotter (1989b)

4.3.1: Experimental Designs

While LRS (1989a) did yield some of the answers to the questions posed above, there are still some questions left unanswered. First, LRS (1989a) leaves unanswered the question of why the behavioral change observed in RS (1989) appears in the first place? In other words, what are subjects trying to achieve by shifting their behavior away from linearity as the prior distribution of the subjects become more skewed? 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, i.e. 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, LRS (1989b) employed an experimental design in which live subject sellers engaged in 75 round sealed bid experiments with pre–programmed computerized buyers (in some cases the computerized subjects were sellers) employing one of two fixed strategies. These functions were the following:

Linear	Broken
$b = .82V$	$b = \begin{cases} V \text{ if } V \leq 60 \\ 60 \text{ if } 60 \leq V \leq 100 \end{cases}$

Clearly by holding the actions of one side of the market constant they are able to more intelligently observe the behavior of the other side and investigate whether subjects are able to infer how their computerized partners are behaving and make an appropriate best response.

In addition to varying the computerized subjects bid function, these experiments were run under two information conditions and using two different sets of priors ($r_1 = 1$, $r_2 = 1$, and $r_1 = .2$, $r_2 = .4$). In the high information condition subjects were told the 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 RS (1989) and LRS (1989a)). The bid functions used were in some sense representative of those we observed in LRS (1989a) while the coefficients used for the prior distributions were the ones which yielded the most dramatic linear and broken behaviors. Clearly, given LRS (1989a) a 75 horizon experiment was also a necessity.

Seventeen experiments were performed 14 of which had the computer playing the role of the buyer and 3 of which had the computer playing the role of the seller. The first eight experiments were paired experiments in which subjects played first against a computerized program using either a linear or broken strategy and then played against a computerized program with the opposite strategy. 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. Clearly, these experiments are of use in answering the second and third questions raised above, since they allow one to measure the extent to which subjects were able to infer what their computerized partners (who were using a stationary strategy) were doing and to measure their ability to choose a best response. Extensive post–experiment surveys were also used. This design does not help in answering the first question posed 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[2] ($r_1 = 0$, $r_2 = 0$). 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 Linhart, Radner and Schotter (1989b) ran two additional experiments. First they ran a certainty experiment ($r_1 = r_2 = 0$) which they hoped to use as a basis of comparison for their previously skewed and uniform distribution experiments (LRS (1989a)). In this experiment it was common knowledge that buyers

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

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 here in the effect of certainty on joint subject behavior.

4.4: Results

The LRS (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 opponent 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 we compare 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 of experiments in which the bidding strategy of the computer is held constant but the prior distributions change does indicate a significant amount of change. Hence, bidding behavior appears to be far more sensitive to the parameters of the mechanism than to the bidding strategy of one's opponent (or at least one's computerized opponent). While this appears strange at first glance, it is not as surprising once one investigates the best response functions of live subjects facing these computerized subjects.[3] For example, Figure 3 below presents the best response 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. To demonstrate how close subjects came to making best responses, we present figure 4 in which the theoretical best response and the average observed bidding functions (solid lines) are plotted on the same graph. Again by average function we mean that function derived by taking the mean slopes and break

3 Peter Linhart furnished the the calculations upon which this statement is made.

Figure 3
Best Response Function v.s. Linear and Broken Computerized Bid Functions

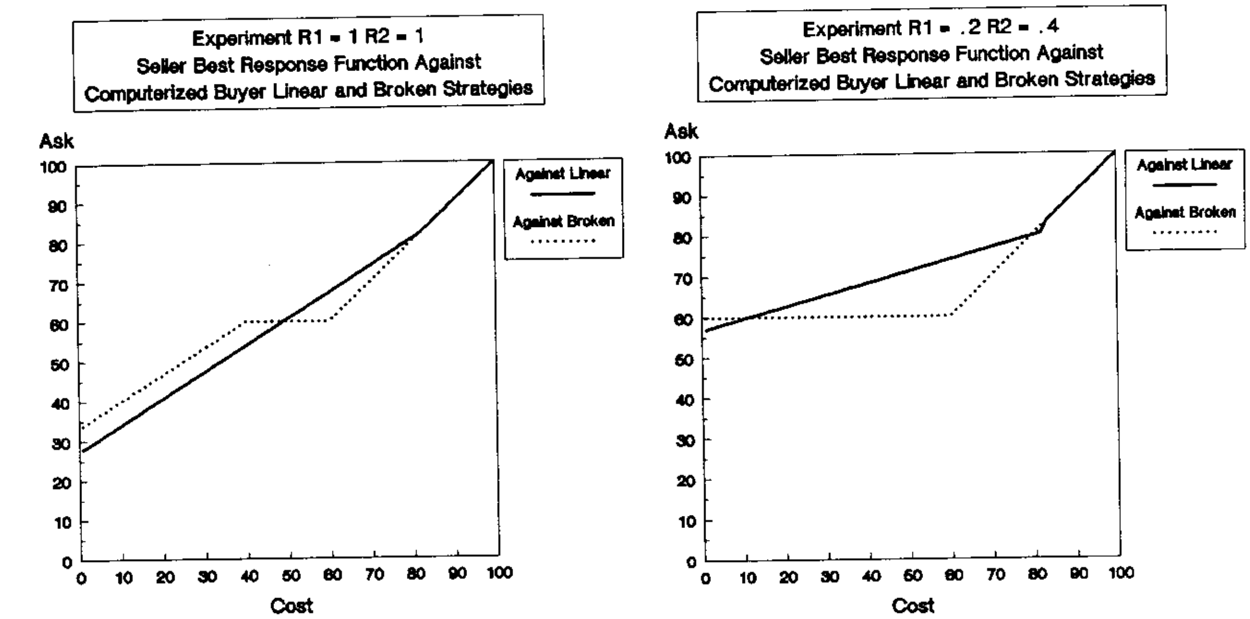

Note: For costs above 60 we have assumed that sellers will bid their costs. Actually any bid of at least their cost will maintain their functions as best responses. Bidding their cost weakly dominates all others, however.

Figure 4
Predicted v.s. Actual Mean Best Response Functions

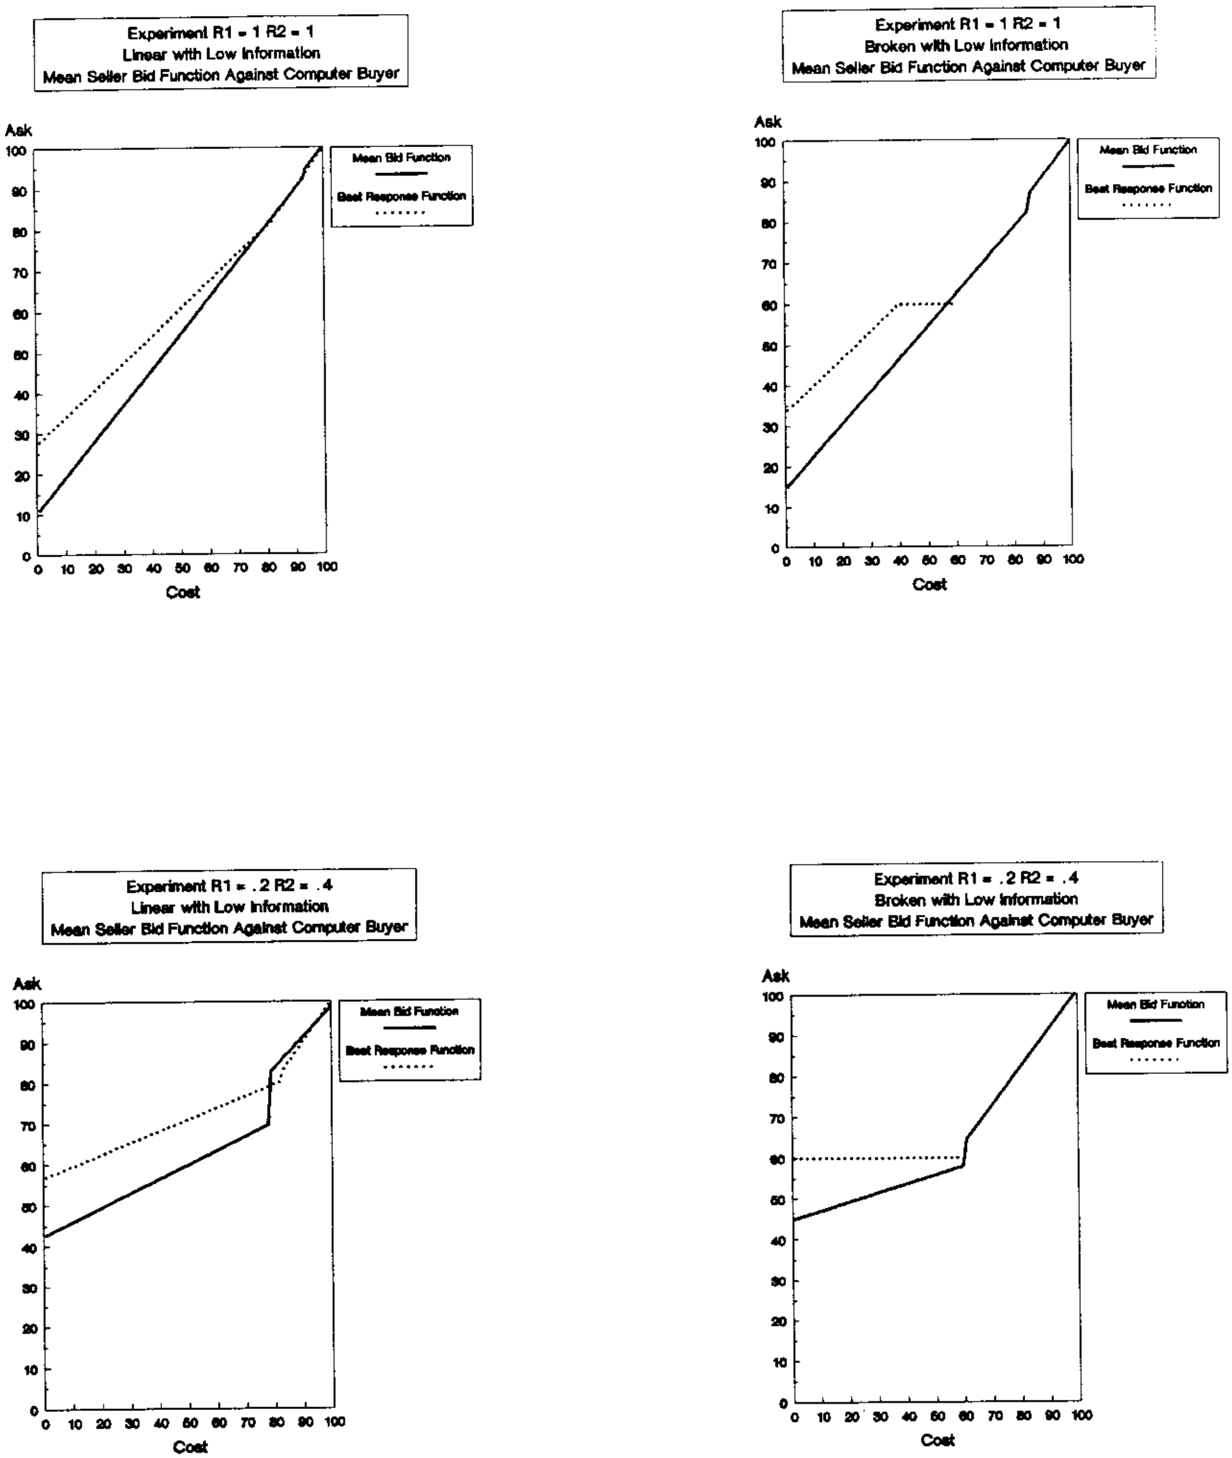

point of the best fitting two–piece piece–wise linear line generated by the scatter of points observed for each subject. Note that by and large the fit appears to be quite good and certainly at least qualitatively, moves in the correct direction.

These graphs indicate also that our subjects were quite adept at inferring what type of bid strategy their computerized partner was using and at taking the appropriate best response to that behavior. Note that since the best response function was not very sensitive to the exact function used by the computer, as long as it is linear or broken, once a subject made the assumption that his computerized partner was using a linear or broken function there was nothing left to learn in the experiment–– the best response became a function of the parameters of the experiment which were common knowledge.

The final conjecture of LRS (1989b) concerned the cause for the shift in bidding behavior as the parameters of the experiment changed. It was conjectured here that in those experiments where the prior distributions were sufficiently skewed, subjects would behave as they would in 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 5, which presents a representative sample of bid and ask time patterns for buyers and sellers, substantiates this belief. 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 which, 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 which is above (below) the agreed upon price. Such behavior would be consistent with our two–piece piece–wise linear function exhibited in Figure 1 where those graphs are drawn in value–bid (cost–ask) space. Figure 6 presents another representative sample of the bid and ask time patterns in the $r_1 = 0.2$, $r_2 = 0.4$ live experiments of LRS (1989a). Note that

Figure 5

Bid - Ask Functions Over Time: Certainty Case ($r_1 = 0, r_2 = 0$)

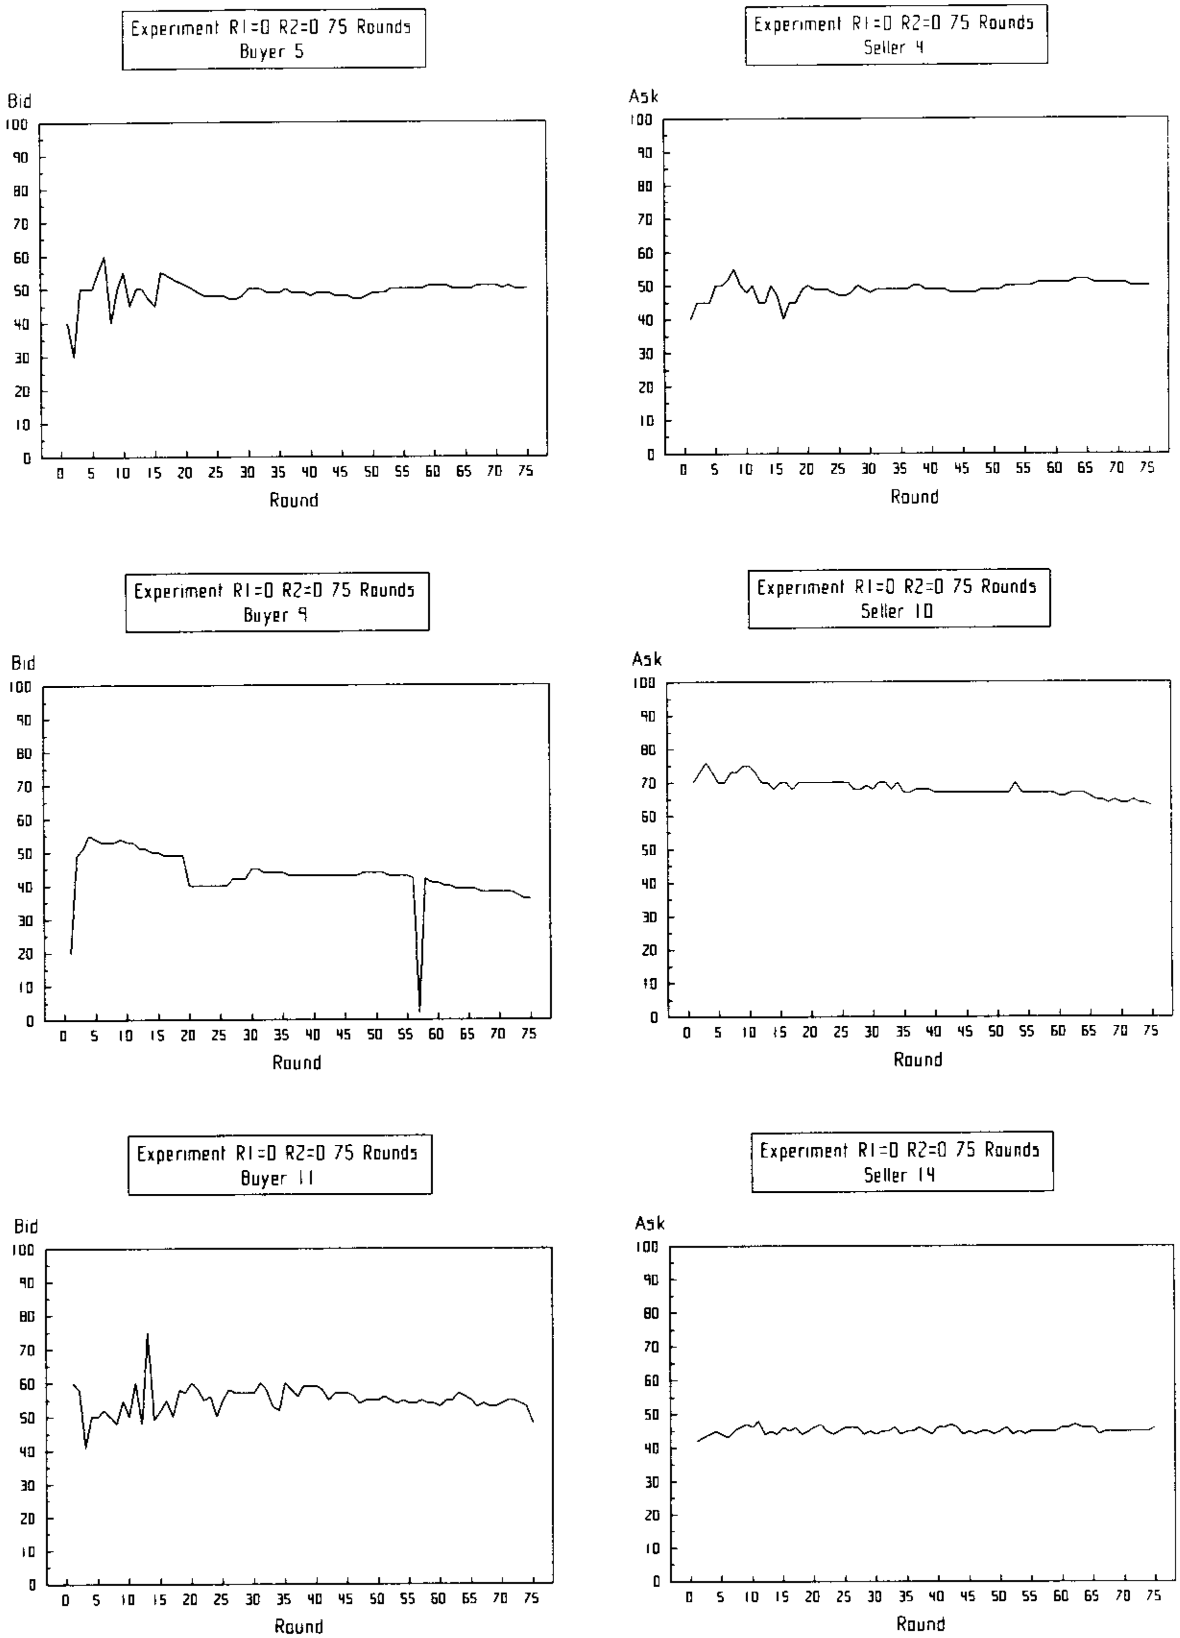

Figure 6

Bid - Ask Functions Over Time: Skewed Distribution ($r_1 = 0.2, r_2 = 0.4$)

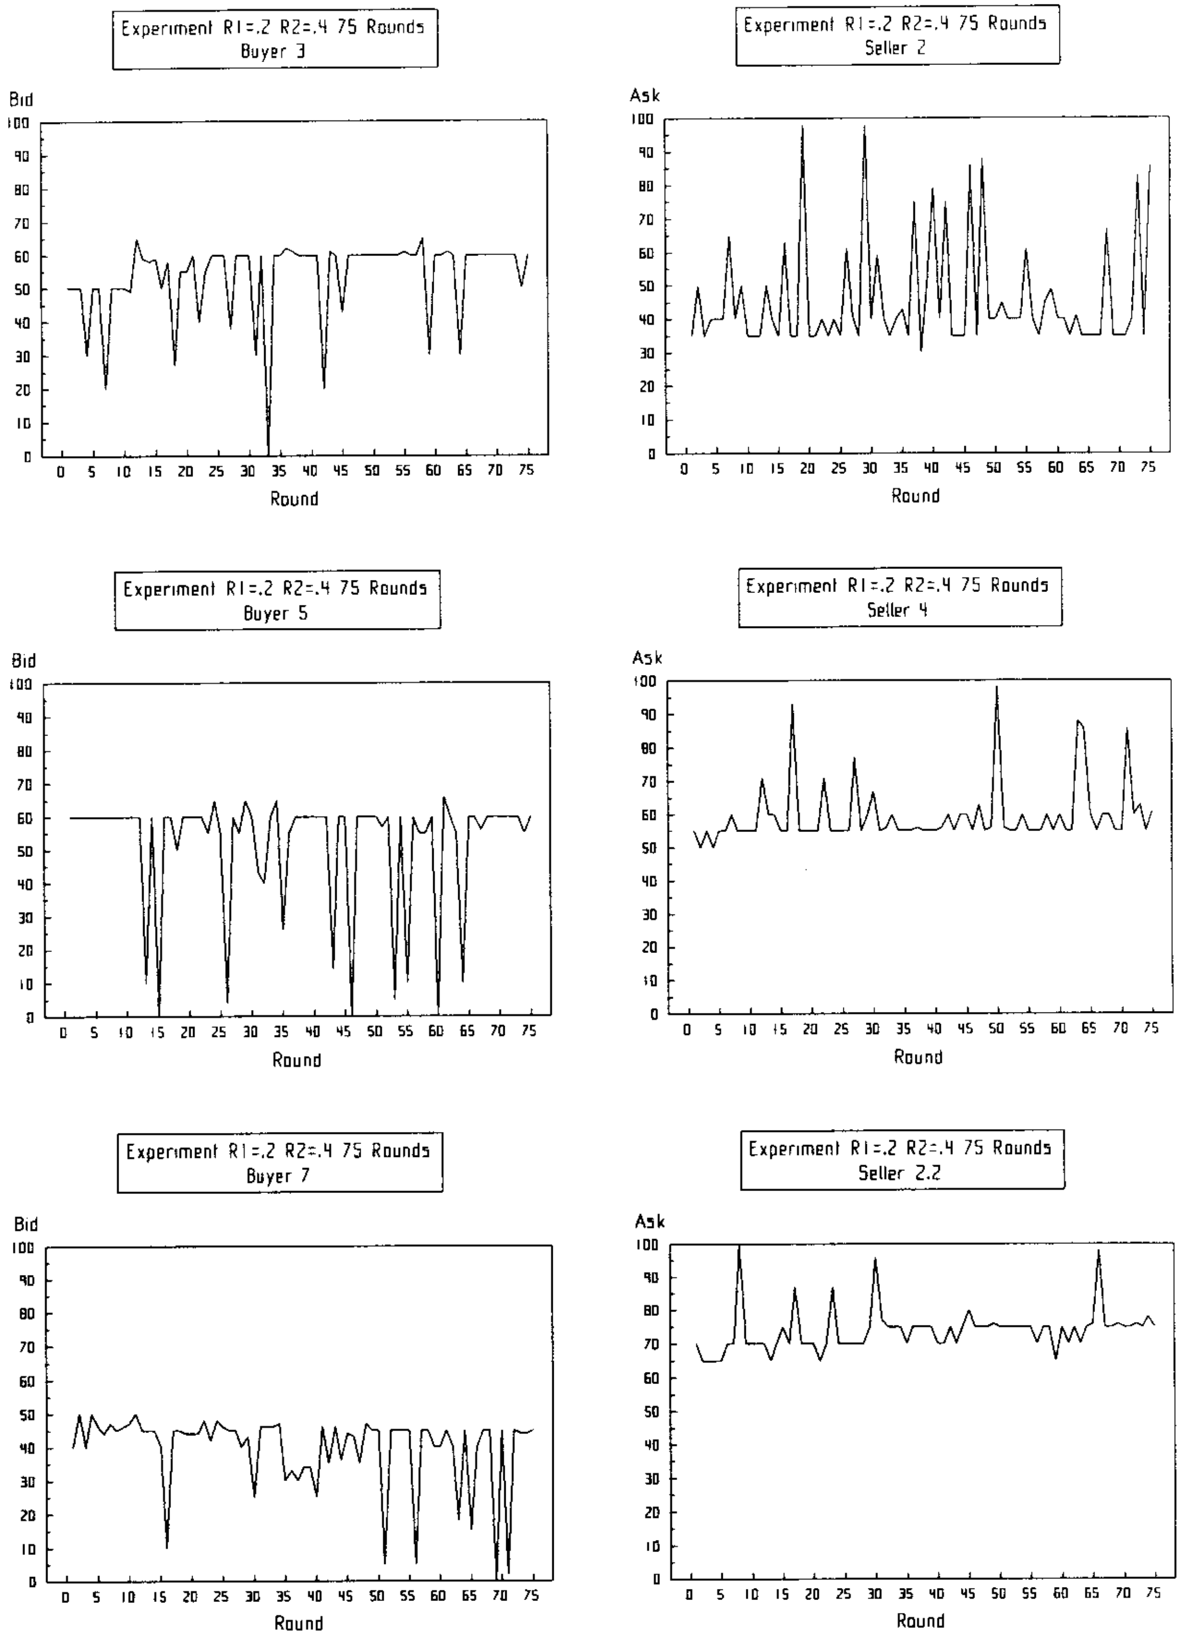

Figure 7
Bid - Ask Functions Over Time: Skewed Distribution Spikes Removed

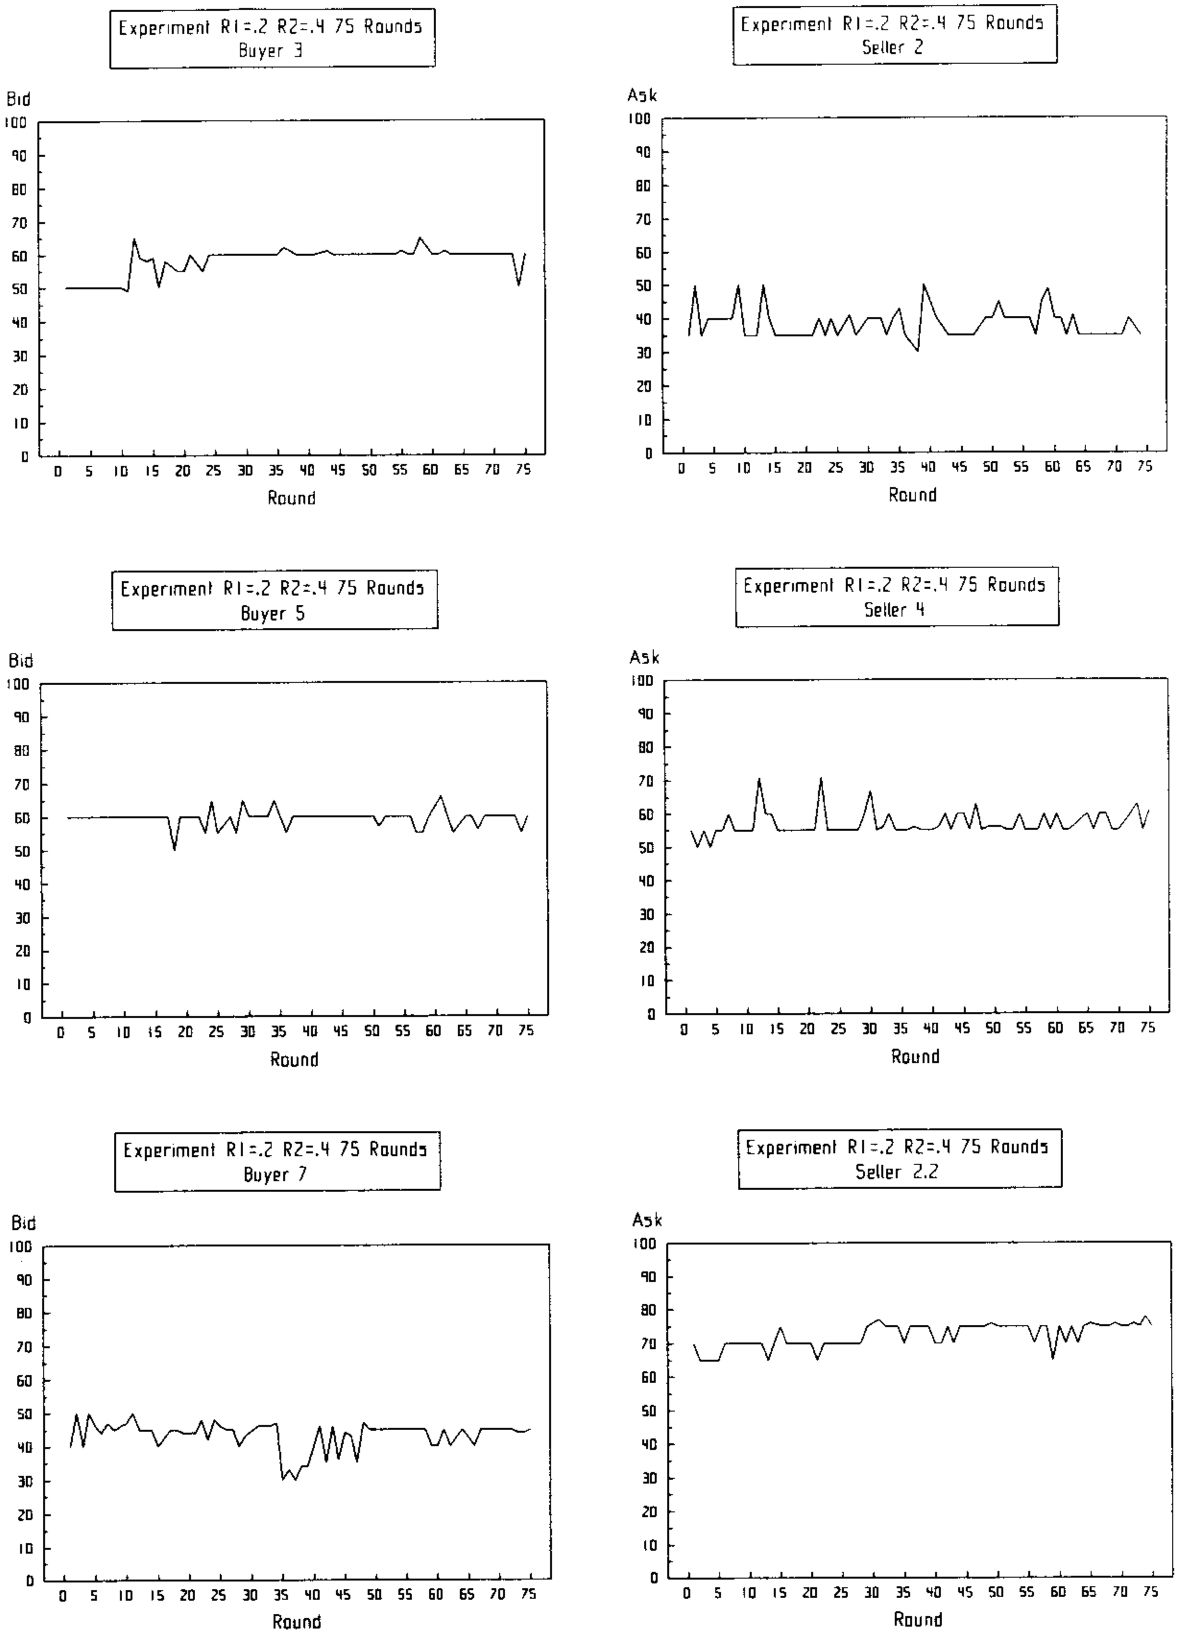

Figure 8
Bid - Ask Functions Over Time: Uniform Distribution ($r_1 = 1, r_2 = 1$)

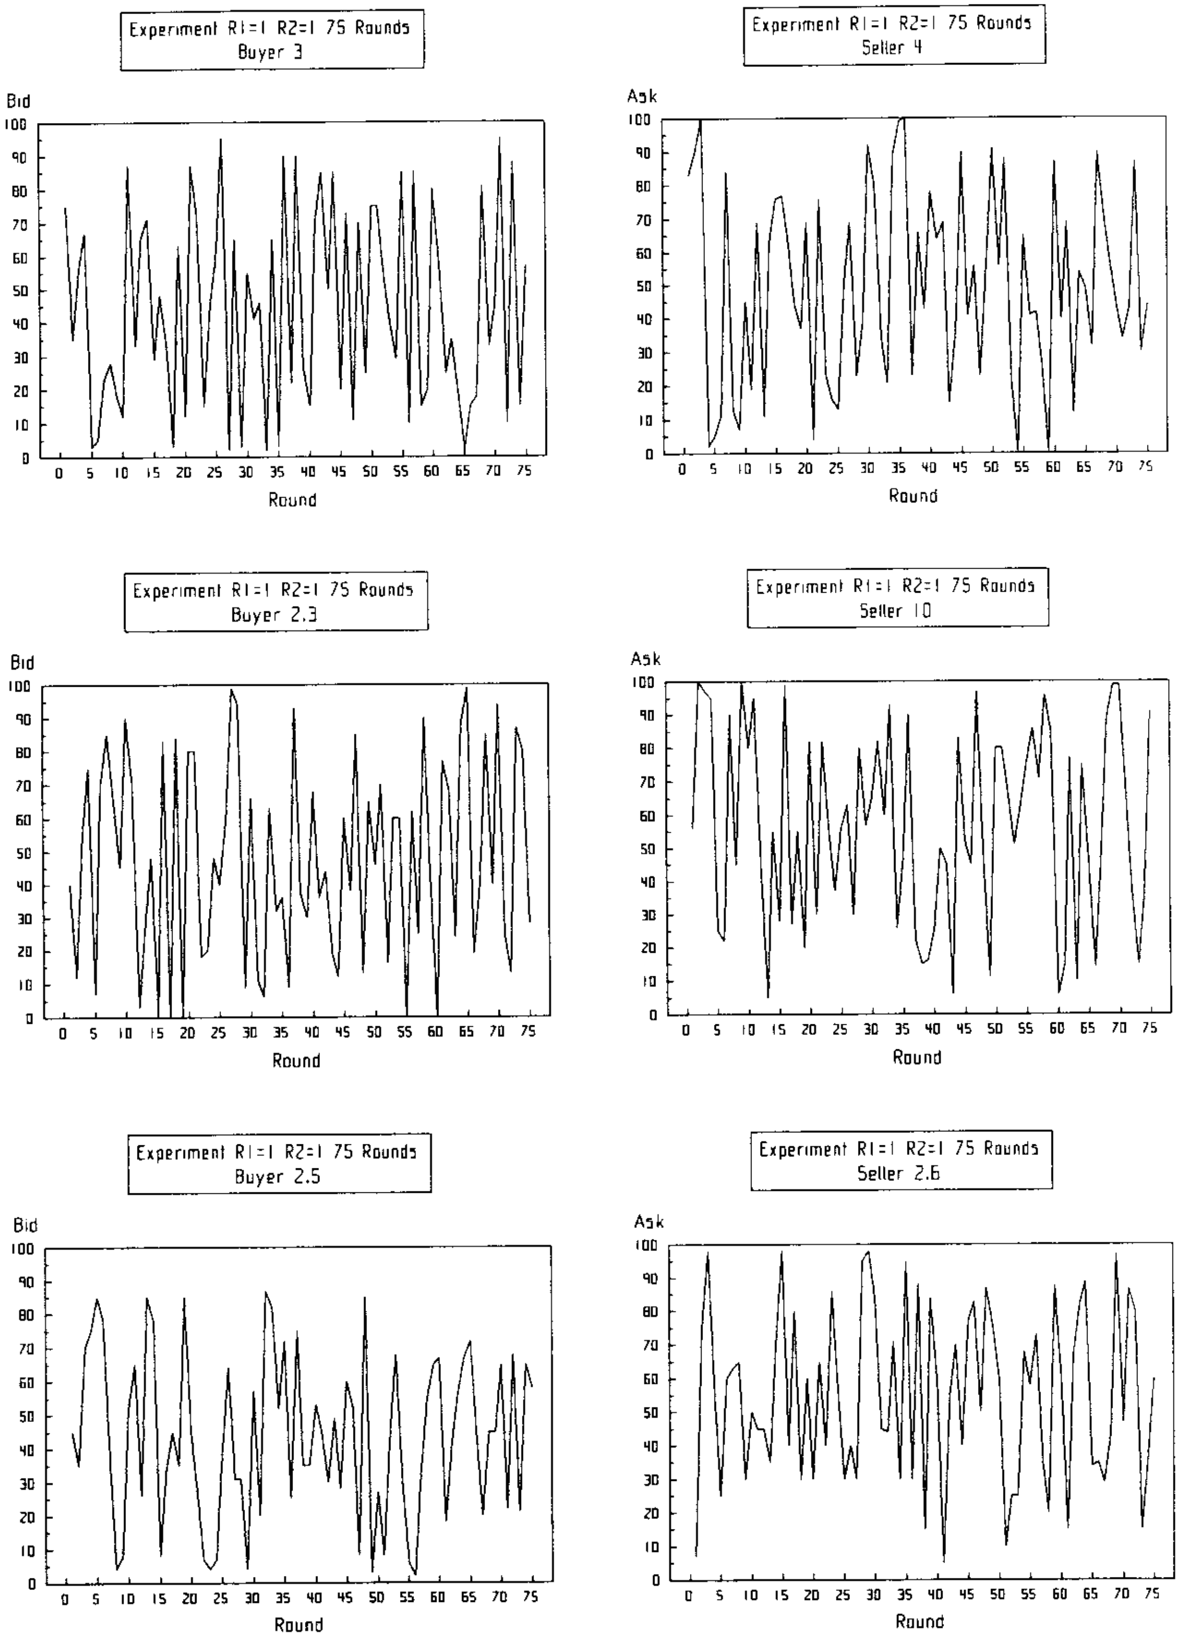

these functions are basically constant with respect to time except for spikes which occur when a value or cost is realized above or below implicitly agreed to price. These spikes are, of course, necessary in order to avoid losses on the parts of subjects. If we were to eliminate these spikes, we would create Figure 7 which has much of the appearance of the graphs in Figure 5. The comparison of Figures 5 and 7 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 8 we present another representative sample of bid patterns from the uniform experiments. Note that this behavior is indeed different from those seen in figures 5, 6, and 7.

Section 5: 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 effect the efficiency of the sealed–bid mechanism.

References

Chattejee, K. and Samuelson, W., "Bargaining Under Incomplete Information", Operations Research, 1983, pp. 835–851.

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

Myerson, R., and Satterthwaite, M., Efficient mechanisms for bilateral trading, J. Econ. Theory, 29 (1983), pp. 265–281.

Leininger, W., Linhart, P., and Radner, R., "The sealed–bid Mechanism for Bargaining With Incomplete Information", Journal of Economic Theory, 48, June 1989, pp. 63–107.

Linhart, P., Radner, R., and Schotter A., "Behavior and Efficiency in the Sealed–Bid Mechanism": Part 1, mimeo, 1989a

Linhart, P., Radner, R., and Schotter A., "Behavior and Efficiency in the Sealed–Bid Mechanism": Part 2", in progress, 1989b

Radner, R., and Schotter A., " The Sealed–Bid Mechanism: An Experimental Study", Journal of Economic Theory, 48, June 1989, pp. 179–221.

Samuelson, W., and Bazerman, M., "The Winner's Curse and Bilateral Negotiations", in, Research in Experimental Economics, Vernon Smith ed., vol.3, JAI Press 1985

Satterthwaite, M., and Williams, S., " Bilateral Trade With the Sealed–Bid k–Double Auction: Existence and Efficiency", Journal of Economic Theory, 48, June 1989, pp. 179–221.

Wilson, R., "Incentive Efficiency of Double Auctions", Econometrica, 53, 1985, pp.1101–1115.