The Economic Journal, 113 (*April*), 477–494. © Royal Economic Society 2003. Published by Blackwell Publishing, 9600 Garsington Road, Oxford OX4 2DQ, UK and 350 Main Street, Malden, MA 02148, USA.

LEARNING TO BID – AN EXPERIMENTAL STUDY OF BID FUNCTION ADJUSTMENTS IN AUCTIONS AND FAIR DIVISION GAMES*

Werner Güth, Radosveta Ivanova-Stenzel, Manfred Königstein and Martin Strobel

We examine learning behaviour in auction and fair division experiments with independent private values under two different price rules, first and second price. Participants play all four games repeatedly and submit complete bid functions rather than single bids. This allows us to study how institutional changes are anticipated and whether learning is influenced by the structural differences between games. We find that learning does not drive bidding towards the benchmark solution. Bid functions are adjusted globally rather than locally. Directional learning theory offers a partial explanation for bid changes. The data support a cognitive approach to learning.

Learning has become a major topic in economic research. Whereas formerly one either relied on (common knowledge of) rationality or that markets would drive out irrational modes of behaviour, one now is interested in the processes of behavioural adjustments and open to what may be their results. In this paper we report an experiment on learning behaviour in four different types of games, the first price (or respectively second price) auction and the first price (or respectively second price) fair division game with independent private values (in all cases). Subjects participated in a sequence of games against randomly matched opponents, with the game type changing after every three rounds of bidding. Every subject was asked for a complete bid function rather than a single bid. A bid function specifies a bid for each possible private value.

The experimental design offers an interesting environment for studying learning in bidding tasks: first, developing a bid function is more demanding than bidding for a single value which is the usual practice in bidding experiments. Second, since the game rules change over time, the task can be considered as rather complex. Adjustment behaviour in complex settings might differ from adjustment behaviour in simple, stationary games. Third, despite differences in game rules the format of decisions and the information feedback on game outcomes are the same across all game types. According to noncognitive learning theories – for example, 'reinforcement learning' (Bush and Mosteller, 1955; Roth and Erev, 1995) – behavioural adjustment should therefore be invariant across game types. However, if learning is driven by cognition and deliberation, bid adjustments should differ between game types.

To distinguish between cognitive and noncognitive learning theories is one way to structure different approaches to learning; see Selten (1998) who distinguishes between 'cognition' as reasoning and 'adaptation' as routine adjustment without

* We gratefully acknowledge helpful comments by David De Meza and two anonymous referees. For financial support we thank the Deutsche Forschungsgemeinschaft (SFB 373) and the EU (TMR research network ENDEAR).

reasoning. Among the cognitive models are, for example, 'best–reply dynamics', 'fictitious play' and 'direction learning'; see, e.g., Selten and Buchta (1998). Among noncognitive models 'reinforcement learning', 'imitation' (Vega-Redondo, 1997), and 'replicator dynamics' (Taylor and Jonker, 1978) are quite well known.

Drawing a borderline between cognitive and noncognitive learning theories is to some extent arbitrary.[1] Even those theories which we consider as noncognitive require some basic reasoning. For instance, to reason that 'past success is a good predictor of future chances' is implicit in reinforcement learning. 'What is good for others is good for oneself' is a reasoning underlying imitation dynamics. However, noncognitive learning theories do not require any understanding of the true game structure in the sense of understanding the links between strategy choices and payoffs. This is different for theories which we call cognitive since they all require a more complete understanding of the game structure. 'Best-reply dynamics' and 'fictitious play' require that the decision maker can determine his payoff for all possible strategy profiles. 'Direction learning' requires this for some, but not necessarily all strategy profiles. We argue for cognitive rather than noncognitive learning and provide experimental evidence supporting this view. However, in spite of the evidence for systematic learning patterns, the data do not suggest where behaviour will converge, if it does so at all. The dispersion in individual behaviour is rather large even when participants are more experienced. Consequently, in real world auctions and fair division games one should not rely on quick convergence to stationary bidding.

The paper proceeds as follows: Section 1 describes the experimental games and their benchmark solutions as well as the experimental procedures and the payments. Section 2 presents our empirical results and Section 3 concludes.

1. Auctions and Fair Division Games

1.1. *Games and Theoretical Solutions*

Bidding behaviour is a favourite topic in experimental economics; for a selective survey, see Kagel (1995). As in auction theory, one distinguishes open and sealed-bid auctions. We will focus here on sealed bid experiments in which a single object is to be allocated and for which each potential buyer has an independent private value. We investigate four different allocation rules referred to as game types (see Table 1): First Price Auction (A1), Second Price Auction (A2), First Price Fair Division Game (F1) and Second Price Fair Division Game (F2). Fair division games differ from auctions with respect to the ownership of the good to be allocated. In F1 and F2 the good is owned by the bidders themselves. The price at which the object is sold is equally distributed among all bidders. In auctions the price is earned by an outside agent, the seller (owner). An example for a fair division game is the allocation of inheritance. The object is collectively owned by the heirs who,

[1] For examples in the psychological literature on learning, see Macias (1996) and Edelmann (1986).

Table 1

The Four Game Types

Price Rule	Auction	Fair Division Game
Price = highest bid	A1	F1
Price = 2nd highest bid	A2	F2

in many cases, are the only bidders. Similar problems result when a joint venture is terminated.[2]

Let v_i be a bidder's private value for the object to be sold, and suppose v_i is drawn for each bidder $i = 1,\ldots,n$ independently from a uniform distribution on the unit interval. If all bidders are risk neutral, the equilibrium bid function $b_i^*(v_i)$, expected equilibrium price $E(p^*)$, and expected equilibrium payoff $E[\pi_i^*(v_i)]$ are as shown in Table 2. For a derivation of these results, see Güth and van Damme (1986). Illustrations of the equilibrium bid functions are provided below together with the presentation of results.

The solutions rely on the assumption that values and bids are continuous variables. In experiments, however, this is usually not the case. In our experiment both, values and bids, are discrete. This raises questions regarding the relevance of the benchmark solutions above.

For a first price auction Goeree *et al.* (1999), using a specific example, have shown that the solution of the continuous case may be a good approximation of the discrete case.[3] However, Riley (1989) shows for the case of discrete values and continuous bids that pure strategy equilibria do not exist. The intuition underlying this result

Table 2

Bid Function, Expected Price, and Expected Payoff According to the Risk Neutral Equilibrium for the Four Game Types

Price	Auction	Fair Division Game
Highest bid	$b_i^*(v_i) = \frac{n-1}{n} v_i$	$b_i^*(v_i) = \frac{n}{n+1} v_i$
	$E(p^*) = \frac{n-1}{n+1}$	$E(p^*) = \left(\frac{n}{n+1}\right)^2$
	$E[\pi_i^*(v_i)] = \frac{v_i^n}{n}$	$E[\pi_i^*(v_i)] = \frac{v_i^n}{n} + \frac{n-1}{n(n+1)}$
2nd highest bid	$b_i^*(v_i) = v_i$	$b_i^*(v_i) = \frac{n}{n+1} v_i + \frac{1}{n+1}$
	$E(p^*) = \frac{n-1}{n+1}$	$E(p^*) = \frac{n^2+1}{(n+1)^2}$
	$E[\pi_i^*(v_i)] = \frac{v_i^n}{n}$	$E[\pi_i^*(v_i)] = \frac{v_i^n}{n} + \frac{n-1}{n(n+1)}$

[2] For an experimental study on a related topic, the so-called zero-revenue auctions, see Franciosi *et al.* (1993). However, since these are multi-unit auctions, we do not compare their results to ours.

[3] The model features two bidders whose values are drawn from {0,2,4,6,8,11}. Since bids had to be integer, this is a rather coarse bid grid.

is as follows: If several bidders exhibit the same value (due to discreteness) and submit the same winning bid, then the winner will be determined randomly (among these bidders). In this case, however, a bidder would have benefited from a marginally higher bid, which destroys the existence of pure strategy equilibria.

Similar arguments hold if the discreteness of bids is sufficiently fine. Thus, pure strategy equilibria may not exist in the games analysed here, either. We are not aware of a theoretical study that derives a mixed-strategy equilibrium for the case of discrete bids and discrete values, neither for first price auctions nor for fair division games. Such an analysis would certainly be beyond the scope of our study and beyond the capabilities of experimental participants. Rather, we will rely here on the continuous case as a reasonable approximation of the experimental situation and use the theoretical solution as a benchmark to be compared with observable bidding.

1.2. *Experimental Games and Procedures*

In our experiment the private values $\tilde{v}_i$ were drawn from the set

$$\tilde{V} = \{50, 60, 70, 80, 90, 100, 110, 120, 130, 140, 150\}$$

with all values $\tilde{v}_i \in \tilde{V}$ being equally likely. These values are denoted in a fictitious currency ECU (Experimental Currency Unit) used by subjects to resell the object to the experimenter. Bids $\tilde{b}_i$ had to be integer values as well, with $\tilde{b}_i$ between 0 and 200 ECU. For ease of comparison of the empirical bids $\tilde{b}_i$ and values $\tilde{v}_i$ with the theoretical solution given above, all our analysis will be done for normalised bids b_i and values v_i:

$$v_i = \frac{\tilde{v}_i - 50}{100}$$
$$b_i = \frac{\tilde{b}_i - 50}{100}.$$

Accordingly, the space of possible values is $V = \{0, 0.1, \ldots, 1\}$. The space of possible bids is $b_i \in \{-0.50, -0.49, \ldots, +1.50\}$. Thus, the experimental rules allow subjects to submit bids below (above) the minimal (maximal) reselling value.

Within a session each subject participated in 36 consecutive games of the four different types. Nine subjects formed a session group. In each of the 36 rounds (periods) they were randomly partitioned into three groups of three bidders. The number of bidders involved in each game ($n = 3$) was known by all, but not their identity. In all sessions all subjects played the same sequence of games: A1 in periods $t = 1$ to 3, A2 in $t = 4$ to 6, F2 in $t = 7$ to 9 and F1 in $t = 9$ to 12. This comprised the first block of 12 periods. Then they played block 2 (periods 13 to 24) and 3 (periods 25 to 36) in the same sequence as block 1. We had no strong reasons for the sequencing A1, A2, F2, F1 within each block. We merely chose this design in order to change only a single institutional aspect (price rule or ownership structure) at a time. Furthermore, we started with the most familiar situation, the first price auction.

Most participants were students of economics or business administration at Humboldt University of Berlin. They had been invited by leaflets to participate in

an experiment announced to last about three hours, and sessions actually took about that long. After entering the laboratory they were placed at isolated computer terminals. Communication among participants was not allowed during the session. While reading the instructions (see Appendix A), they could privately ask for clarification or for help in handling the PC.

In each game they had to submit a complete bidding strategy (bid vector) $\tilde{b}_i(\tilde{v}_i)$. Thus, they had to enter a bid for each of the 11 values $\tilde{v}_i \in \tilde{V}$. The actual value $\tilde{v}_i'$ was drawn thereafter. Payments were determined according to the game rules and the submitted bidding strategies.[4] Each participant was informed about $\tilde{v}_i'$, about whether or not he was buyer, about the price $\tilde{p}$ at which the object was sold, and about their own payoff in that game. Then the next game followed. Appendix B shows some sample screen shots of the computer software.

Thus, each game type was played nine times. In the first of these nine games the bid fields were blank, and each subject had to enter a vector of 11 bids (one for each $\tilde{v}_i \in \tilde{V}$). In later periods the last bid vector for the same game type was displayed as default. It could be revised or submitted as it is. Of course, this may favour the status quo and may work against adjustments of behaviour over time. We did it for practical reasons. If subjects did not want to always adjust all bids, this would save time and help to prevent boredom by the task. Altogether we ran six sessions and collected 1,944 bidding strategies (54 subjects times 36 games).

1.3. *Payments*

Subjects total earnings out of the 36 games ranged between DEM 31 and DEM 96 with a mean of DEM 56 (about US$ 33 at the time of the experiment including a show up fee of DEM 10). In the first half of the sessions we used the same conversion rate for ECU into cash: 1 ECU = DEM 0.05. Theoretically and practically this generates rather asymmetric monetary incentives for auctions versus fair division games. Güth (1998) tried to guarantee equal monetary incentives by adjusting the conversion rate such that equilibrium profits were equal for $v_i = 0.5$. Instead, we used actually observed profits of the first three sessions and adjusted the conversion rate to induce equal expected payoffs based on observed behaviour. This meant for sessions 4 to 6 that one ECU was worth DEM 0.2857 in auctions and DEM 0.02857 in fair division games.

Essentially this meant that we had a payoff treatment: three sessions with equal conversion rate and three sessions with unequal conversion rate. Theoretically these payoff differences are irrelevant. Since in all data analyses we ran, we did not find them relevant, we will not discuss them any further.

[4] The strategy method obviously provides more information than collecting only one bid for a single value. But since ex post only one component of the bid vector is payoff relevant, it lowers the incentives of bidding at each single value. By restricting the set V, we tried to achieve a reasonable compromise between information and incentives.

2. Results

2.1. *Description of Raw Data*

Figure 1 shows the distributions of individual bids for each game type (A1 to F2) and different levels of experience (block 1 to 3). Each dot represents a participant's average bid for the given reselling value (i.e., the average of the three bids of subject i in the respective block). The top row displays the bids for A1. The reference lines indicate 'true value'-bidding (black line) and benchmark bidding (grey line). Bidding above the benchmark for the game A1 is rational for risk averse bidders; namely, this may increase the probability of winning, at the cost of earning less in case of winning. But in A1 bidding above the true value can not be rationalised, since this is (weakly) dominated. Thus, the two lines define an area for rational bidding in case of risk aversion. In fact, we find many observations inside this area.[5] But there is also a substantial number of observations outside this area. Overbidding the true value or underbidding the A1-benchmark occurs more often in block 1 than with experience (blocks 2 and 3).

Rows 2, 3, and 4 present the data for A2, F1, and F2, respectively. Again, the black lines indicates true value bidding, the grey line indicates benchmark bidding (in case of A2 the two lines coincide). While many theoretical models in auction theory work with the assumption that bidders are risk neutral (see, e.g., the survey by Wolfstetter (1996)), we find overall that the benchmark solutions based on the continuous case for risk neutral bidders cannot explain the data very well. In A1 and A2 as well as in F1 we observe a strong tendency of overbidding the benchmark.[6] In F2 risk aversion calls for bidding below the benchmark. In contrast, Figure 1 shows a large number of bids above the grey reference line. We think that participants perceive F2 as the most complex game among the four experimental games, since fair division games seem less familiar than auctions and since the second price rule seems less familiar than the first price rule. However, similar to A1, A2, and F1 even in F2 the number of bids that are inconsistent with risk aversion is reduced with experience. This indicates some kind of bid adjustments at the aggregate level. In the following sections we will investigate these adjustment processes at the individual level.

In a companion paper (Güth *et al.*, 2002) we analyse static rather than dynamic aspects of behaviour. For instance, we compare prices and efficiency rates across game types. We show that most bid functions are increasing and (almost) linear. Finally, even though bidding differs from the benchmark solution, it is in line with comparative statics prediction (across game types).[7]

[5] Further experimental evidence on risk averse bidding in first price auctions is provided, e.g., by Cox *et al.* (1982, 1985, 1992). The December 1992 issue of AER contains a controversial debate on risk aversion as an explanation for bidding behaviour.

[6] Such overbidding in first price single-unit private value auctions has been found as well, e.g., by Cox *et al.* (1988) and by Dyer *et al.* (1989). In second price sealed-bid auctions, bidding above the dominant strategy was also observed, e.g., by Kagel and Levin (1993).

[7] For related experimental confirmation of such comparative statics predictions, see Kagel and Levin (1993).

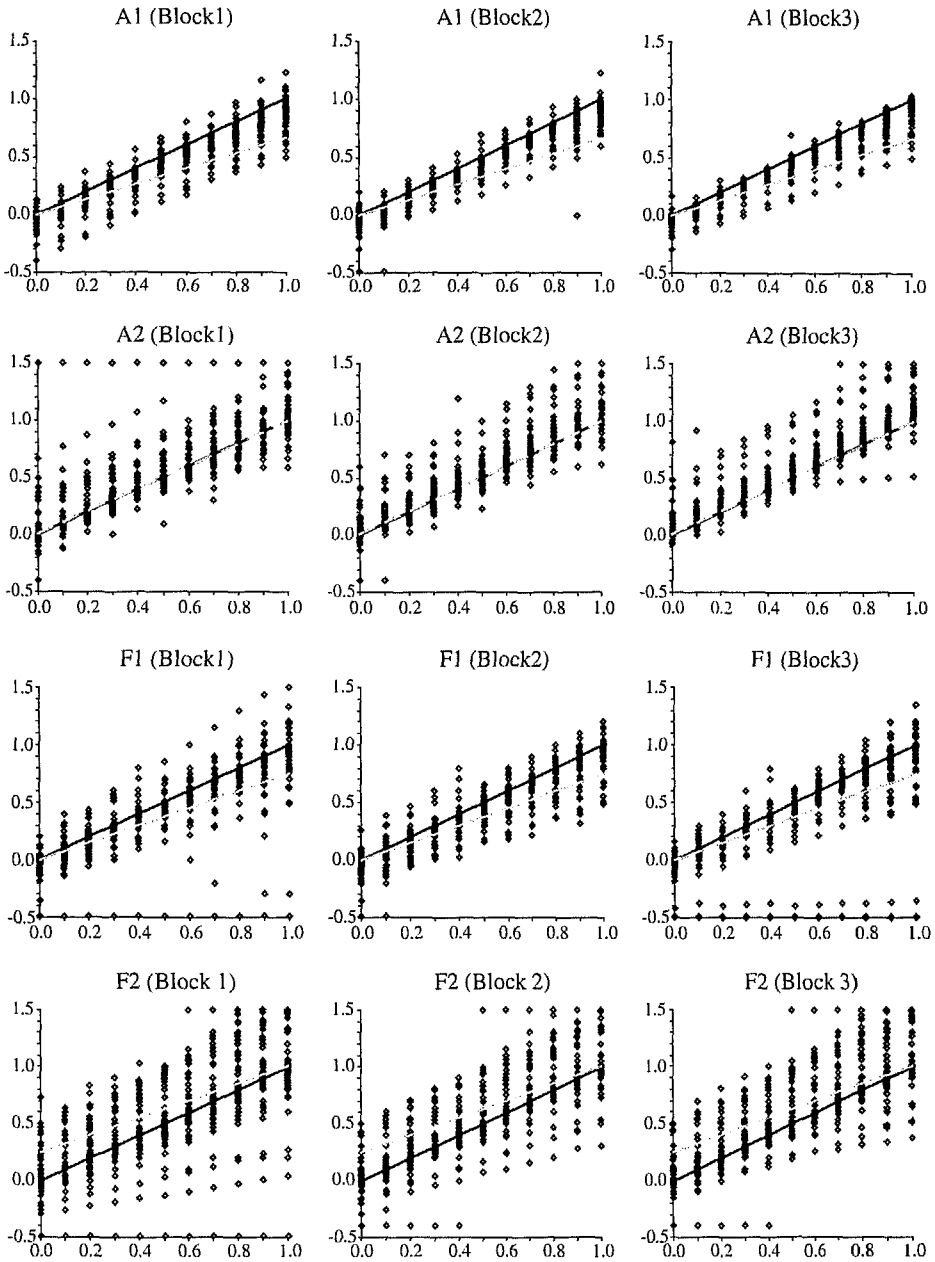

Fig. 1. *Scatterplot of the Average Bids for Each Subject, for Three Levels of Experience, and for All Game Types (the horizontal axes denote the private values, the vertical axes denote the bids)*

2.2. *Monotonicity and Convergence of Bid Changes*

Before asking whether subjects learn a specific kind of behaviour, one might ask whether they do learn anything at all. If so, one should observe more bid function

adjustments early on and some stationary behavioural pattern after sufficient time. Thus, with experience bidding should converge toward some stable individual bid function. To investigate this remember that each subject i played each type of game nine times. For every type we will refer to the bid function in the ninth play as i's final bid function. Hence the final bid functions of game types A1, A2, F2, and F1 are $b_{i,27}(v_i)$, $b_{i,30}(v_i)$, $b_{i,33}(v_i)$, and $b_{i,36}(v_i)$ respectively. To measure bid function adjustment we use the Euclidean distance between i's bid function $b_{it}(v_i)$ in period t and i's final bid function. For instance, in A1 we calculated

$$D_i^{A1}(t) \equiv \left\| b_{it}(v_i) - b_{i,27}(v_i) \right\|$$

for all subjects i and all periods $t < 27$ in which A1 was played, i.e. $t \in \{1, 2, 3, 13, 14, 15, 25, 26\}$. Analogously, we calculated $D_i^{A2}(t)$, $D_i^{F2}(t)$, and $D_i^{F1}(t)$. We consider an adjustment process as monotone if $D_i^j(t)$ with $j \in \{A1, A2, F2, F1\}$ is decreasing in t. Furthermore, a monotone adjustment process will be called 'convergent' if $D_i^j(t)$ decreases more rapidly in earlier than in later periods, that is, if $D_i^j(t)$ is convex. Both, monotone processes and (even more so) converging processes, will be interpreted as evidence for learning.

For classification of the observed processes we used slightly weaker criteria than those described above in order to allow for some error. Specifically, we fitted a piecewise–linear regression line to the data with $D_i^j(t)$ as dependent and t as independent variable allowing for a kink of this line after four (out of the eight) periods. Accordingly, a process is regarded monotone if both slope coefficients are negative. If, in addition, the coefficient is smaller in absolute value for later periods, the process is convergent.

Table 3 displays the relative frequencies of observed adjustment processes. As one can see, between 61% and 74% of all individual bid functions exhibit a convergent adjustment process; between 92% and 100% are monotone. Thus, even though we do not know yet what specific kind of bidding behaviour subjects learn, we do know that they learn something.

2.3. *No Learning of Bidding According to the Benchmark Solution*

As shown the data are not in line with the theoretical benchmark. On the other side, when looking more thoroughly at Figure 1, one can recognise some changes in the bidding behaviour as subjects gain experience. In all four game types

Table 3

Classification of Adjustment Processes for Different Game Types (%)

	Game Type			
Adjustment Process	A1	A2	F2	F1
Monotone	92	95	96	100
Convergent	61	69	74	70
Not convergent	31	26	22	30
Not monotone	8	5	4	0

subjects learned to avoid irrational bidding. For instance, in A1 there is a tremendous reduction of bids above the 'true value' bidding as well as of extreme overbidding and/or underbidding in A2 and F1. Since there is a learning process going on, it could well be that this process works into the direction of the benchmark solutions. Figure 1 has indicated already that there is no strong convergence towards them. However, it is interesting to investigate this in more detail. One might speculate whether subjects will play according to the benchmark solutions if they are given more time to gain experience and to learn. Figure 2 shows the development of bidding for each session. Specifically, based on piecewise-linear regression estimates of the aggregate bid functions (of each session), it presents the time paths of predicted bids for $v_i = 0.5$ and each game type.[8] A reference line in each figure indicates the benchmark solution b_i^*. There is no obvious movement towards it. We observe similar patterns for all other reselling values $v_i \neq 0.5$. Therefore, we conclude that the learning process does not drive towards the benchmark solution.

2.4. *Local Versus Global Changes*

The information feedback received by subject i at the end of each game – the actually drawn value $\tilde{v}_i'$, the price $\tilde{p}$, whether or not i bought the object, and i's profit – may suggest whether $b_i(v_i')$ should be adjusted in future periods. But this feedback does not tell anything regarding bids $b_i(v_i'')$ for all values $v_i'' \neq v_i'$. This seems obvious for a noncognitive learner. But if learning is cognitive, it is conceivable that changing $b_i(v_i')$ induces the bidder to reconsider his bid function at other values v_i'' as well. A naturally arising question is therefore whether bid functions were adjusted only 'locally' at v_i or at many values. If they were adjusted at all 11 values, we will call this a 'global' adjustment.

Figure 3 shows that most adjustments are in fact global. It reports frequency distributions of bid changes for each game type. For all four game types the mode of the respective distribution is at 11; that is, in these cases subjects adjusted their bid function at each value. Furthermore, a more detailed look at the data reveals that, conditional on whether the bid function was changed at all, it was changed globally in 54% of the cases in A1. The according statistics for A2, F2, and F1 are 48%, 34%, and 33%, respectively. Thus, if a bid function is changed at all, changing all bids is most frequent. Bid functions are adjusted globally rather than locally, suggesting a more cognitive type of learning.

Global adjustments are examined further in Table 4. It reports the number of monotone shifts versus non-monotone shifts. Monotone shifts are such that for each value the respective bid is increased (upward shift) or decreased (downward

[8] To give the benchmark a better chance we excluded some data which were obviously problematic. Remember that altogether we collected 1944 individual bid functions. Güth *et al.* (2002) showed that 97% of these were strictly increasing and that most of them were quite accurately predicted (according to the coefficient of determination R^2) by a piecewise-linear model – i.e., $b_i(v_i)$ is piecewise-linear in v_i – allowing for a kink of the regression line at $v_i = 0.5$. In computing the estimates of the aggregate session bid function, which we present here, we therefore excluded individual bid functions which were not strictly increasing, and furthermore those with an R^2 of the piecewise-linear model smaller than 80%. Taken together, 84 (4%) of the bid functions were excluded.

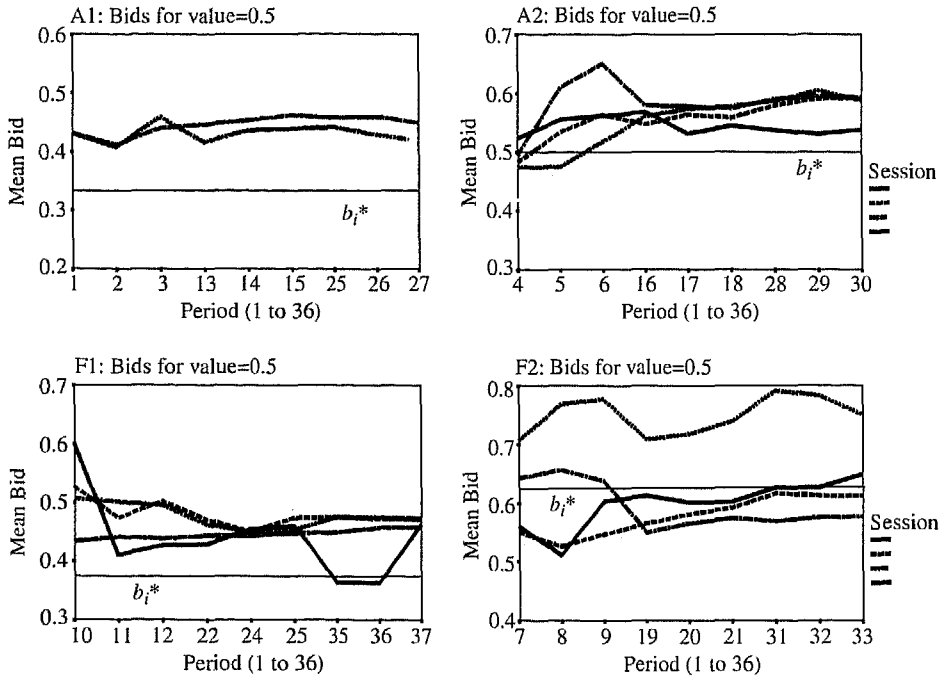

Fig. 2. *Time Paths of Estimated Bids for $v_i = 0.5$ and All Game Types*

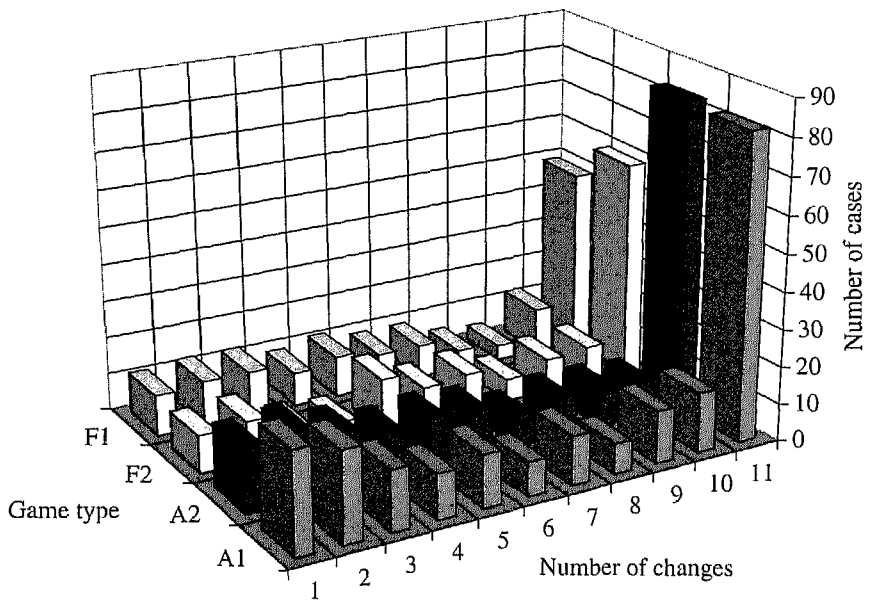

Fig. 3. *Frequency Distributions for the Number of Changes within the Bid Function for the Different Game Types*

Table 4

Frequencies of Different Kinds of Global Adjustments

	Game Type			
Type of Global Adjustment	A1	A2	F2	F1
Monotone shift	73	77	55	44
Parallel shift up	19	33	20	14
Other shift up	19	23	19	12
Parallel shift down	16	9	5	7
Other shift down	19	12	11	11
Non–monotone shift	9	6	4	5
Total	82	83	59	49

shift). In many cases all bids were adjusted by the same amount. This is denoted as a parallel upward or, respectively, downward shift. Monotone shifts and, even more so, parallel shifts are highly structured forms of global adjustments. Both findings, first, that local feedback leads to global adjustments in many cases and, second, that adjustments are highly structured, support a cognitive learning approach.

2.5. *Direction Learning*

We now investigate 'direction learning', a learning process proposed by Selten and Buchta (1998) which predicts the direction of change of a strategy.[9] It is based on an *ex post* comparison between the success of the chosen strategy and some alternative strategy. Specifically, direction learning proposes that, if a strategy is adjusted, it will be adjusted into the direction where an improvement would have been feasible. Direction learning requires the decision maker to determine the payoff he would have received in period $t-1$ according to a fictitious strategy profile.

In a first price auction direction learning proposes that a subject i who became a buyer in period $t-1$ should lower his bid in period t since a bid reduction might have increased his payoff in $t-1$. A bid increase in $t-1$ might have done nothing but decrease his payoff so that an improvement was feasible only in one direction. Similarly, a non-buyer j whose value in $t-1$ was above the price ($v_j^t > p$) should increase his bid in t since this might have led to a positive profit in $t-1$. However, a non-buyer j whose value in $t-1$ was below the price ($v_j^t \leq p$) could not possibly have made a positive profit, either by increasing or by decreasing his bid in $t-1$. Hence, direction learning theory does not make a prediction in this case (Selten and Buchta, 1998).

By similar arguments we determined directional predictions for all four game types. They are summarised in Table 5. Compared to Selten and Buchta, we applied somewhat stronger criteria. To see this consider, for example, a buyer i in A1 who earned a positive profit in $t-1$. A bid reduction might have

[9] See also Güth (1998).

Table 5

Conditions for Applying Direction Learning Theory and Predicted Directions of Bid Changes

Predicted direction of bid change	Criterion for directional prediction			
	A1	A2	F2	F1
Bid ↓	Buyer and $\Pi_i \leq 0$	Buyer and $\Pi_i \leq 0$	Buyer and $\Pi_i < \frac{1}{3}p$	Buyer and $\Pi_i < \frac{1}{3}p$
Bid ↑	Not buyer and $p < v_i^l$	Not buyer and $p < v_i^l$	Not buyer and $p < v_i^l$	Not buyer and $p < v_i^l$
No prediction	Otherwise	Otherwise	Otherwise	Otherwise

increased his profit, but he might as well have lost money by becoming a non-buyer (since bids had to be made in discrete steps, this could have happened even for marginally lower bids). Selten and Buchta (1998, p. 88) argue that 'it is not clear how high the second bid was but a higher profit could have been made by some lower bid'. This statement is acceptable for continuous bids. But with discrete bids, even a small bid reduction may lead to i becoming a non-buyer. Thus, it is not clear whether an improvement is feasible. Consequently, direction learning may not apply. However, if buyer i earned a non-positive profit in $t-1$, a bid reduction could have never reduced, but may have increased his profit. This is the condition we applied in A1. Analogous reasoning holds for the other game types.

According to these criteria, direction learning is applicable in 318 of 1,728 cases (≈18%).[10] Table 6 reports frequencies and percentages of correct or, respectively, incorrect predictions. Considering all four game types (see the rightmost column), 131 out of 318 cases (41%) are in line with direction learning and only 12 cases (4%) are not. In 175 cases (55%) no bid change occurred. Thus, if a bid change occurred at all, it was correctly predicted by the theory in the majority of cases.

Table 6

Predictive Success of Direction Learning Overall and for Each Game Type. Absolute and Relative Predictions (Changes)

Observed direction	A1	A2	F2	F1	All
Right	34	39	33	25	131
	(56%)	(57%)	(34%)	(27%)	(41%)
Wrong	2	1	5	4	12
	(3%)	(1%)	(5%)	(5%)	(4%)
No bid change	25	29	59	62	175
	(41%)	(42%)	(61%)	(68%)	(55%)
Total	61	69	97	91	318
	(100%)	(100%)	(100%)	(100%)	(100%)

[10] Note that directional learning does not make a prediction for first round decisions.

The conditional frequency for adjusting into the right direction is 92%. Comparing across game types, the percentage of correct predictions is smaller in fair division games (34% and 27% respectively) than in auctions (57% and 56% respectively). We attribute this to the more complex structure of fair division games. Apparently subjects had more difficulties in figuring out the correct direction in more complex environments and thus preferred not to adjust their behaviour. Thus, overall direction learning theory explains a substantial proportion of bid adjustments.

Investigating the data in more detail, we find that the frequency of correct predictions is higher than the frequency of incorrect predictions in each of the six sessions. Accordingly a Binomial test rejects the null hypothesis ('the probability for correct predictions and the probability for wrong predictions are equal') in favour of direction learning theory ($p < 0.016$, $N = 6$ one-tailed test). Furthermore, in all sessions the relative frequency of correct predictions is higher in auctions (A1 and A2 together) than in fair division games (F1 and F2 together). Thus, we conclude that direction learning theory is more successful in explaining adjustments in less complex games ($p < 0.016$, $N = 6$, one-tailed Binomial test).

Finally, Table 7 reports the relative frequency of correct predictions of direction learning theory in periods with recent feedback versus periods with distant feedback. The numbers in parentheses show the number of correct predictions divided by the total number of cases in which direction learning theory is applicable. Recent feedback means that a subject plays the same game as in the previous period. Distant feedback means that a subject returns to game type i after playing nine rounds of other game types j. In these cases one can imagine that direction learning predicts less well than in periods with recent feedback.

According to Table 7 the relative frequency of correct predictions is higher with recent feedback than with distant feedback. This holds overall (see bottom row) and for each game type. However, looking at each session separately, there are four sessions where direction learning predicts better for recent feedback but two sessions in which it predicts better for distant feedback. Thus, we cannot reject the null hypothesis ('predictive success of direction learning is equal for both feedback conditions').

Table 7

Relative Frequencies of Correct Predictions of Direction Learning Theory for Recent versus Distant Feedback

	Game type				
Feedback time	A1	A2	F2	F1	All
Distant	50% (5/10)	55% (11/20)	33% (7/21)	15% (4/27)	35% (27/78)
Recent	57% (29/51)	57% (28/49)	34% (26/76)	33% (21/64)	43% (104/240)

We summarise our findings regarding direction learning as follows:

1. About 18% of the observed bids can be considered as test cases for direction learning theory.
2. The frequency of bid adjustments that are correctly predicted by direction learning theory is significantly higher than the frequency of wrong predictions.
3. Direction learning theory predicts significantly better for auctions than fair division games; that is, for the less complex game types.
4. Direction learning theory predicts slightly better for recent feedback than for more distant feedback, but the differences are not statistically significant.

Since direction learning requires a cognitive model of a payoff function that determines the payoff for fictitious strategy profiles, we regard it as a cognitive theory. According to Selten (1998) direction learning is not 'a full-fledged learning theory, but rather the description of an important aspect of behaviour' (p. 423). 'It is a qualitative theory which makes only weak predictions' (p. 422). In general, it does not specify by which amount behaviour is adjusted and whether the process converges. Nevertheless, direction learning 'describes an influence of cognition on adaptation' (Selten, 1998, p. 422) and may organise empirical data.[11]

3. Conclusions

In this paper, we study learning behaviour in a laboratory experiment on four different game types. So far, there 'has been very little study of learning and adjustment processes in private value auctions' (Kagel, 1995, p. 521) and fair division games. A few exceptions are, for example, Smith and Walker (1993), Selten and Buchta (1998) and Güth (1998). Compared to some other experimental studies on learning,[12] the decision environment in our experiment is more complex. The game rules differ between game types, although the format of decisions remains constant. The participants submit bid functions rather than bids for a single private value. We classify different modes of adjustment behaviour into cognitive versus noncognitive (adaptive) learning models and characterise some basic aspects of learning behaviour which, in our view, suggest cognitive rather than noncognitive learning.

Without claiming to know where learning will finally settle, we find the following features of the learning process: First, we observe no strong tendency to learn bidding according to the benchmark solution for the continuous case (bids and values) and risk neutral bidders. While this model may be viewed as a natural candidate to approach the bidding problem theoretically, it does a bad job in explaining behaviour. Second, we show, however, that subjects do learn something: Almost all adjustment paths are monotone, and most are convergent

[11] For other experimental contexts, see, e.g., Cason and Friedman (1997), Nagel (1996).

[12] For instance, Daniel *et al.* (1998) study learning in a bilateral bargaining situation. Abbink *et al.* (2001) investigate learning within the ultimatum game.

in the sense that adjustments are smaller in later rather than in earlier periods. Third, bid functions are adjusted globally rather than locally. Thus, subjects 'interpret' local feedback as having informational content for other or even all components of the bid vector. A cognitive principle that accounts for such behaviour is 'generalisation'.[13] Fourth, global adjustments are highly structured in the sense that most are monotone shifts, and many are even parallel shifts. Last but not least, direction learning offers a partial explanation of the observed changes in bidding behaviour.

Despite these systematic learning patterns, the data do not suggest where behaviour will converge to, if it does so at all. The dispersion in individual behaviour is rather large even with experience. Given these findings it seems unlikely that behaviour in real world auctions and fair division games will converge quickly to stationary bid functions. In order to draw a convincing picture of learning behaviour, more work needs to be done, including not only theoretical modelling, but also collecting more data and formulating stylised facts.

Max Planck Institute for Research into Economic Systems, Jena
Humboldt University of Berlin
University of Erfurt
International Institute of Infonomics, Heerlen

Date of receipt of first submission: September 1999
Date of receipt of final typescript: May 2002

Appendix: A. Instruction[14]

Please read these instructions carefully. They are identical for all participants.

During the experiment you will take part in several auctions. In every auction a fictitious commodity is for sale which you can resell to the experimenters. You are one of three bidders. Each bidder has his own private reselling value v which can be 50, 60, 70, 80, 90, 100, 110, 120, 130, 140, or 150 ECU (Experimental Currency Unit) and is independently drawn. Each value appears with the same probability.

Before you learn your individual reselling value v, you have to place a bid for every possible v:

$$b(50), b(60), b(70), \ldots, b(150).$$

After every bidder in your group has placed his bid vector your actual bid is determined by $b(v)$. The bidder with the highest bid buys the commodity and pays a price according to the pricing rule. Then he sells the commodity to the experimenter and receives his reselling value. The other bidders do not pay anything and do not receive the commodity. If there are two or three highest bids, the buyer is chosen at random.

[13] Noncognitive approaches could possibly model this behaviour as parameter learning, e.g., by allowing adjustments of a proportional or absolute degree of under- or, respectively, overbidding that is applied to all components of the bid vector.

[14] This is a shortened and translated version of the instructions. For the original instructions (in German), please contact one of the authors.

There are four different types of auctions. In type 1 and 2 the auction revenue is kept back by the experimenter whereas in type 3 and 4 the auction revenue is equally divided among the bidders. In auction types 1 and 4 the price corresponds to the highest actual bid. In auction types 2 and 3 the price which has to be paid corresponds to the second highest actual bid.

Type 1 (*First Price Auction*)

- Price = highest bid ($p = b_1$)
- Bidder with highest bid becomes buyer. He pays p.
- Revenue (p) is kept back by the experimenter.
- Profit of buyer: $v - p = v - b_1$
- Profit of non-buyers: 0

Type 2 (*Second Price Auction*)

- Price = second highest bid ($p = b_2$)
- Bidder with highest bid becomes buyer. He pays p.
- Revenue (p) is kept back by the experimenter.
- Profit of buyer: $v - p = v - b_2$
- Profit of non-buyers: 0

Type 3 (*Second Price Fair Division Game*)

- Price = second highest bid ($p = b_2$)
- Bidder with highest bid becomes buyer. He pays p.
- Revenue (p) is distributed among the bidders.
- Profit of buyer: $v - p + \frac{1}{3}p = v - b_2 + \frac{1}{3}b_2 = v - \frac{2}{3}b_2$
- Profit of non-buyers: $\frac{1}{3}p = \frac{1}{3}b_2$

Type 4 (*First Price Fair Division Game*)

- Price = highest bid ($p = b_1$)
- Bidder with highest bid becomes buyer. He pays p.
- Revenue (p) is distributed among the bidders.
- Profit of buyer: $v - p + \frac{1}{3}p = v - b_1 + \frac{1}{3}b_1 = v - \frac{2}{3}b_1$
- Profit of non-buyers: $\frac{1}{3}p = \frac{1}{3}b_1$

Altogether you play 36 auctions. In each auction the bidder groups are formed randomly. After you have placed your bid you are informed about the price, your private reselling value, whether or not you bought, and how much you have earned. Any decision you make is anonymous and cannot be related to you. If you have questions, please, do not ask aloud, but raise your hand. We will then clarify problems privately.

B. Sample screen shots

The screen in Figure B1 was used by the subjects to place their bids. In the upper right corner one finds information about the type of the game. When subjects play a certain type of game for the first time all bid fields are empty. In subsequent rounds subjects look at the strategy they used in the last play of the same game type. They do not have to retype their strategy if they do not want to change it.

Figure B2 is the screen which subjects received after an auction. It informs the subject whether he became the buyer, the price, the individual value v_i^*, his own bid for this value, and the payoff resulting from all these events.

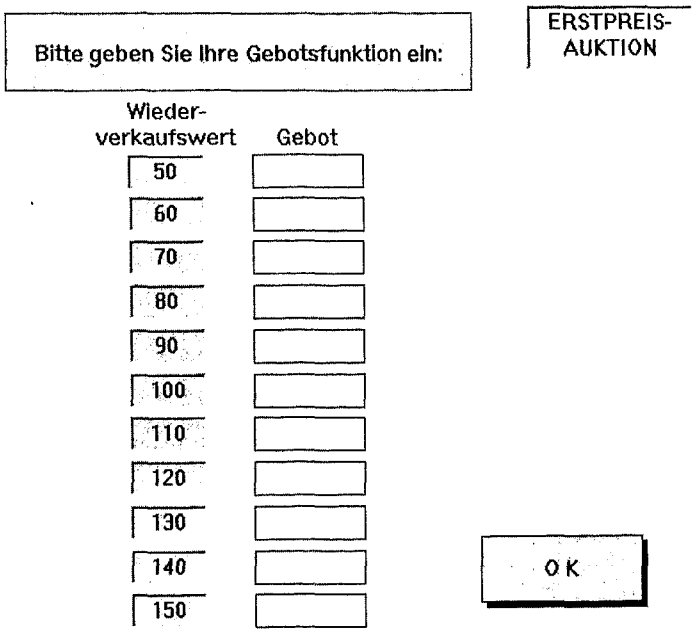

Fig. B1. *Example of a Bidding Screen*

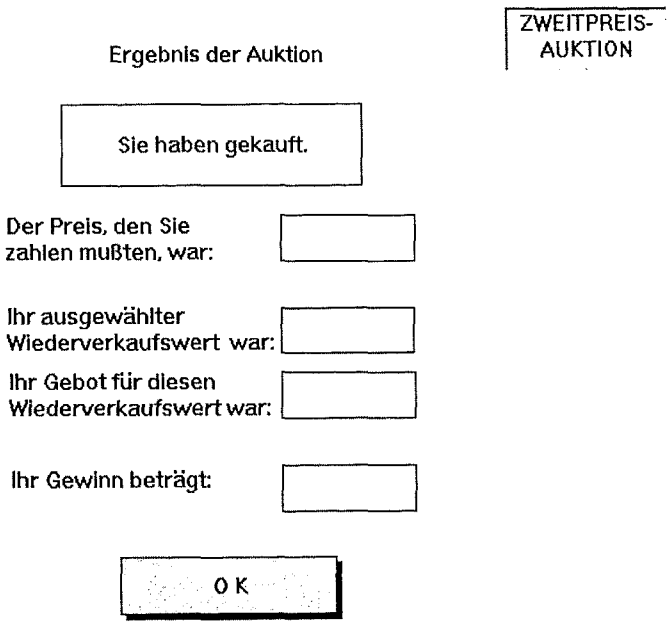

Fig. B2. *Example of an Information Screen*

References

Abbink, K., Bolton, G., Sadrieh, A. and Tang, F. (2001). 'Adaptive learning versus punishment in ultimatum bargaining', Games and Economic Behaviour, vol. 37, pp. 1–25.

Bush, R. and Mosteller F. (1955). *Stochastic Models for Learning*, New York: Wiley.

Cason, T. N. and Friedman, D. (1997). 'Price formation in single call markets', *Econometrica*, vol. 65 (March), pp. 311–45.

Cox, J. C., Roberson, B. and Smith, V. L. (1982). 'Theory and behavior of single object auctions', in (V. L. Smith ed.), *Research in Experimental Economics*, vol. 2, pp. 1–43. Greenwich, Conn.: JAI Press.

Cox, J. C., Smith, V. L. and Walker, J. M. (1985). 'Experimental development of sealed-bid auction theory: calibrating controls for risk aversion', *American Economic Review*, Papers and Proceedings, vol. 75, pp. 160–5.

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

Cox J. C., Smith, V. L. and Walker, J. M. (1992). 'Theory and misbehavior of first-price auctions', *American Economic Review*, vol. 82, no. 5, pp. 1374–443.

Daniel, T. E., Seale, D. A. and Rapoport, A. (1998). 'Strategic play and adaptive learning in the sealed-bid bargaining mechanism', *Journal of Mathematical Psychology*, vol. 42, pp. 133–66.

Dyer D., Kagel, J. H. and Levin, D. (1989). 'Resolving uncertainty about the number of bidders in independent private-value auctions: an experimental analysis', *Rand Journal of Economics*, vol. 20, pp. 268–79.

Edelmann, W. (1986). *Lernpsychologie*. München: Psychologie Verlags Union Urban & Schwarzenberg.

Franciosi, R., Isaac, R. M., Pingry, D. E. and Reynolds, S. S. (1993). 'An experimental investigation of the Hahn-Noll revenue neutral auction for emissions licences', *Journal of Environmental Economics and Management*, vol. 24, pp. 1–24.

Goeree, J. K., Holt, C. A. and Palfrey, T. R. (1999). 'Quantal response equilibrium and overbidding in private-value auctions', *Journal of Economic Theory*, vol. 104(1), pp. 247–72.

Güth, W. (1998). 'On the effects of the pricing rule in auctions and fair division games – An experimental study', in (D. Budescu, I. Erev, R. Zwick, eds.), *Games and Human Behaviour*, Mahwah N.J.: Lawrence Erlbaum Ass., pp. 237–58.

Güth, W., and Damme, E. van (1986), 'A comparison of pricing rules for auctions and fair division games', *Social Choice and Welfare*, vol. 3, pp. 177–98.

Güth, W., Ivanova-Stenzel, R., Königstein, M. and Strobel, M. (2002). 'Bid functions in auctions and fair division games: experimental evidence', *German Economic Review*, vol. 3(4), pp. 461–88.

Kagel, J. H. (1995). 'Auctions: a survey of experimental research', in (J. Kagel and A.E. Roth, eds.), *The Handbook of Experimental Economics* pp. 501–85, Princeton N.J.: Princeton University Press.

Kagel, J. H. and Levin, D. (1993). 'Independent private value auctions: bidding behavior in first-, second-, and third-price auctions with varying numbers of bidders', ECONOMIC JOURNAL, vol. 103, pp. 838–79.

Macias, S. (1996). 'Learning defined', in (F.N. Magill, ed.), *International Encyclopedia of Psychology*, vol. 2. London: Fitzroy Dearborn.

Nagel, R. (1996). 'Unravelling in guessing games: an experimental study', *American Economic Review*, vol. 85, pp. 1313–26.

Riley, J.G. (1989). 'Expected revenue from open and sealed bid auctions', *Journal of Economic Perspectives*, vol. 3, pp. 41–50.

Roth, A.E. and Erev, I. (1995). 'Learning in extensive-form games: experimental data and simple dynamic models in the intermediate term', *Games and Economic behaviour*, vol. 8, pp. 164–212.

Selten, R. (1998). 'Features of experimentally observed bounded rationality', *European Economic Review*, vol. 42, nos. 3–5, pp. 413–36.

Selten, R. and Buchta, J. (1998). 'Experimental sealed bid first price auctions with directly observed bid functions', in (D. Budescu, I. Erev, R. Zwick, eds.), *Games and Human Behaviour*, Mahwah, N.J.: Lawrence Erlbaum Ass., pp. 79–102.

Smith, V. L. and Walker, J. M. (1993). 'Reward, experience and decision costs in first price auctions', *Economic Inquiry*, vol. 31, pp. 237–44.

Taylor, P. and Jonker, L. (1978). 'Evolutionary stable strategies and game dynamics', *Mathematical Biosciences*, vol. 40, pp. 145–56.

Vega-Redondo, F. (1997). 'The evolution of Walrasian behavior', *Econometrica*, vol. 65, pp. 375–84.

Wolfstetter, E. (1996). 'Auctions: an introduction', *Journal of Economic Surveys*, vol. 10 (4), pp. 367–420.