GAMES AND ECONOMIC BEHAVIOR **10**, 39–64 (1995)

Auctions with Artificial Adaptive Agents*

JAMES ANDREONI

Department of Economics, University of Wisconsin—Madison, Wisconsin 53706

AND

JOHN H. MILLER

Department of Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213; and Santa Fe Institute, Santa Fe, New Mexico 87501

Experiments on auctions find that subjects make systematic bidding errors that cannot be explained within the context of Nash equilibrium bidding models. Experimenters and others have conjectured that learning by subjects could lead to errors consistent with those observed. Here, we create and analyze a model of adaptive learning and demonstrate that such a model can capture the bidding patterns evident among human subjects in experimental auctions. Moreover, our model provides a variety of insights into the nature of learning across different auction institutions. *Journal of Economic Literature* Classification Numbers: C7, C9.

1. INTRODUCTION

In recent years there has been a growing interest in the theory of auctions, and in experimental examinations of these theories. An important and challenging result from these experiments is that subjects do not generally adopt Nash equilibrium behavior, but instead make systematic

* The authors are grateful to Ronald Harstad, Donald Hausch, John Kagel, John Rust, Larry Samuelson, Tom Palfrey and two anonymous referees for helpful comments. This research was partially supported by grants to the Santa Fe Institute (including core funding from the MacArthur Foundation, National Science Foundation, and the U.S. Department of Energy). The first author acknowledges support from the National Science Foundation and the Alfred P. Sloan Foundation, and the second author thanks Sun Microsystems. Errors are the responsibility of the authors.

0899-8256/95 $12.00

errors in their bidding strategies.[1] Both experimenters and theorists have puzzled over these "bidder errors," and it is generally perceived that the "misbehavior" of subjects must be addressed from a behavioral point of view if we are to fully understand the experiments and their implications for actual auction behavior.[2]

This paper examines these bidding errors in the context of adaptive learning, and our analysis helps to reconcile the existing theoretical and experimental results. The notion that learning could explain the observed errors is not new. For instance, Kagel and Dyer (1988, p. 196) state that subjects "display clear patterns consistent with genuinely evolutionary learning processes."[3] Friedman (1992) echoed the view that adaptation might usefully explain bidder choices. While the ability to compare experimental outcomes to both rational and adaptive theoretical results has obvious scientific advantages, the difficulty confronting researchers is how to create and analyze general theoretical models of adaptive behavior. Here, we perform numerical analyses on a general class of adaptive learning algorithms.

We study learning by performing "computational experiments" on systems of artificial adaptive agents (AAA). Such computational experiments allow one to study the complex, dynamic, and stochastic processes inherent in these models of learning. The AAA methodology examines the interaction of computational agents modeled by adaptive learning algorithms. The analysis of such systems provides a new approach to understanding fundamental economic and social phenomena that complements existing theoretical and experimental methods (Holland and Miller, 1991). Existing studies using this and related methodologies include Axelrod (1987), Miller (1986, 1995), Marimon *et al.* (1990), Rust *et al.* (1992, 1994), Boylon (1990), Binmore and Samuelson (1992), Crawford (1991), Miller and Andreoni (1991), and Kollman *et al.* (1992).

The AAA systems we study here represent simple learning machines. Unlike real subjects, they are incapable of learning by employing complex counterfactual scenarios or deep introspection, but instead must rely on a simple adaptive search for finding better solutions based on past experience. Nonetheless, such systems give us a lower-bound on the potential impact of learning behavior on such systems, and thus represent an important step in developing an understanding of adaptation in economic games.

[1] See Cox *et al.* (1982, 1983, 1985), Kagel *et al.* (1987, 1988), Kagel and Levin (1985, 1986), Kagel and Dyer (1988), and Dyer *et al.* (1989).

[2] See Harrison (1989) and the recent exchange in the *American Economic Review* by Friedman (1992), Kagel and Roth (1992), Cox *et al.* (1992), and Merlo and Schotter (1992).

[3] Also see Kagel and Levin (1986, p. 917), Kagel *et al.* (1988), and Dyer *et al.* (1989).

We find that adaptive learning is consistent with the main qualitative results from experiments. In addition we find that the various auction institutions display very different adaptive dynamics. Nash equilibrium bidding appears to be easiest to learn in affiliated-values auctions and most difficult to learn in independent-values auctions. While bidders in independent-values auctions have the most complete information on the distribution of the private values of other bidders, our results indicate that this information serves mainly to complicate the task of the bidders. We also find that risk aversion does not yield behavior more consistent with Nash bidding in first-price auctions, which coincides with experimental findings. Also, by comparing small and large groups in common-values auctions, we conclude that the winner's curse observed in large groups may have more to do with the fact that a given error is more costly in a large group than with difficulties in learning or with meaningful deviations from rationality. These and other results suggest that our general approach may be usefully employed to shape future theoretical and experimental studies in this area.

The next section discusses the particular learning algorithm used in this paper. Section 3 describes the structures of the auctions studied. Sections 4 and 5 discuss the results, and section 6 provides a summary and conclusion.

2. An Adaptive Learning Model

This paper models adaptive learning behavior through the use of a genetic algorithm. Genetic algorithms were developed by Holland (1975) as robust methods for adaptive search, learning and optimization in complex problem domains. They are based on natural models of selection and evolution. In these models interacting strategies form a population and compete against one another. Over time strategies that perform poorly are removed, while strategies that perform well are retained and modified. New strategies are created by combining structures from existing strategies through the application of "genetic operators."[4] Such a model has natural analogs to adaptive learning, where agents imitate those strategies that perform well and innovate new strategies by recombining parts of existing strategies.

The algorithm used here has the following design. A population of 40 binary strings (strategies) is randomly generated. Each string encodes two

[4] For a detailed introduction to genetic algorithms see Goldberg (1989).

parameters that define a linear bidding function.[5] Depending on group size (either four or eight), each string participates in a round of auctions (either 100 or 200 auctions respectively[6]). The payoff to each string is the sum of its profits across all auctions conducted during a given round. A new population of 40 strings is then formed by first assigning each string a probability weight based on its relative payoffs in the last round, with higher payoffs yielding greater weight.[7] Using these weights, 20 pairs of strings are randomly drawn, with replacement, from the population. With a 50% probability, a given pair of strings is inserted unaltered into the new population. Otherwise, the pair is subjected to the crossover and mutation operators. Crossover recombines parts of two existing strategies at a randomly chosen point. Mutation alters the state of a single bit of the string. Each bit of the newly crossed strings is mutated with an independent probability of 8% (this rate is exponentially decayed with a half-life of 250 generations). This completes a generation. The new population of strings is again introduced into the auction environment, and the procedure is repeated for 1000 generations. The basic mechanisms of the algorithm model two types of learning. The direct reproduction by performance of the strategies captures the notion of imitating successful strategies. Crossover and mutation allow the innovation of new strategies via the recombination of useful parts of old strategies. Crossover and mutation are simple mechanisms for agents to modify their existing strategies based on past experience. Crossover allows an agent to create a new strategy by borrowing parts of previously successful strategies. Mutation allows an agent to make small modifications to old strategies in hopes of finding something better. Although these operators are simple in action, they have proven to be powerful ways to search nonlinear problems for better solutions (see Goldberg, 1989).

Note that genetic algorithms represent a robust and broad class of adaptive algorithms. Such algorithms only require populations of "solutions" to be reproduced by performance and to have new members created via genetic operators. The algorithms are robust to actual algorithmic and

[5] The first parameter was restricted to the range [0, 2] and the second one to [−2, 2]. Each parameter was coded by 10 bits. The 10 bits were interpreted as a binary integer (if the parameter was signed, then the initial bit determined the sign) and normalized to the appropriate range.

[6] The different number of auctions was designed to give identical information flows under the two group sizes. Recall that each simulation has 40 total bidders. Hence, in the 4-bidder groups there were (100 auctions)/(group) × (10 groups) = 1000 auctions per generation. Similarly, in the 8-bidder groups there were (200 auctions)/(group) × (5 groups) = 1000 auctions per generation.

[7] Strings with payoffs below 1.5 standard deviations from the mean were removed from the population. After these strings were removed the scores were normalized. See Miller (1995) for more detail.

parametric choices (see, for example, Schaffer *et al.* 1989). To verify this we considered a number of variations of both the parameters and algorithm reported above, and found that our results are robust to reasonable algorithmic and parametric changes that do not diminish the level of feedback to the players. For instance, changing the mutation rate has little effect. However, reducing the number of auctions in a round without increasing the number of generations does change the dynamics. In particular, since nonzero payoffs are only received when an auction is won, reducing the number of auctions slows the convergence, although the relative performance of the model across auctions is not affected. The results are also not sensitive to our choice of the domain for the possible parameter values, and increasing this domain in any direction generates similar paths for the system. Thus, the behavior we observe appears to be contained in a much larger equivalence class of adaptive behavior.

3. Theoretical and Experimental Background

In order to gain insights into the behavior of experimental subjects, our computer auctions are deliberately tailored after existing auction experiments. All of the auctions have common parameters and structures in order to facilitate comparisons among them. Let n be the number of bidders. For each individual auction, x_0 and ε are drawn randomly from uniform distributions on the intervals [1000, 2000] and [0, 500] respectively. We then use these parameters to randomly draw n numbers, x_i, from a uniform distribution on the interval $[x_0 - \varepsilon, x_0 + \varepsilon]$. By using these numbers and altering their definitions to bidders, we can construct affiliated and independent private-values and common-values auctions. In this section we review current theoretical and experimental findings concerning these three auction types.

3.1. Affiliated Private-Values Auctions

In private-values auctions, bidders have their own values for the item being auctioned. These private values are affiliated if the bidders do not know the distribution from which all other private-values were drawn, but know only that the higher their own value, the higher others' values are likely to be. This will be true, for example, if an item is purchased for its speculative value. This type of informational assumption is discussed by Milgrom and Weber (1982a).

Our auctions derive their structure from the formulation of Kagel *et al.* (1987). Here, x_i represents bidder i's actual private value for the good. To make the auctions affiliated, bidders are not given full information about the distribution of private values. Hence, in addition to their private value, bidders are only allowed to know ε. For first-price auctions, Kagel *et al.* (1987) show that the Nash equilibrium bidding function in this model is [8]

$$b(x_i) = x_i - \frac{2\varepsilon}{n}. \tag{1}$$

Expected profits of the winning bidder are

$$E\pi = \frac{2\varepsilon}{n}.$$

We also conduct second-price (Vickrey) auctions, in which the object goes to the highest bidder at a price equal to the second highest bid. As is well known, it is a dominant strategy for all bidders to bid their true values. In this case, the predicted Nash equilibrium profits of the winning bidder are the expected difference between the nth and $(n-1)$th order statistic:

$$E\pi = \frac{2\varepsilon}{n+1} \tag{2}$$

Kagel *et al.* (1987) found that subjects in first-price affiliated private-values auctions generally bid in excess of the risk-neutral Nash equilibrium prediction, and made profits that were below the predicted levels. Behavior often differed significantly from the Nash prediction. Nonetheless, they found that the risk-neutral Nash model did a better job of explaining the data than a model of (constant relative) risk aversion, or two ad hoc models.

Kagel *et al.* (1987) also studied second-price auctions. The surprising result from their experiments is that subjects bid well in excess of the dominant strategy. In 80% of the auctions, the market clearing price exceeded the Nash level by significant amounts.[9]

[8] By restricting analysis to signals in the interval $[1000 + \varepsilon, 2000 - \varepsilon]$, they show that the Nash equilibrium bid function is $b(x_i) = x_i - 2\varepsilon/n + y(x_i)/n$, where $y(x_i) = [2\varepsilon/(n + 1)]\exp[-(n/2\varepsilon)(x_i - 1000 - \varepsilon)]$. However, because of its negative exponent, $y(x_i)$, moves negligibly close to zero as x_i increases beyond $1000 + \varepsilon$. We do not believe that this simplification of the bid function has any significant effect on our results or conclusions. This will be discussed in subsequent footnotes.

[9] This result was in clear contrast to the ascending-clock auctions that Kagel *et al.* (1987) designed to test English open-outcry auction, which are, in theory, equivalent to second-price auctions. In this case, bidders almost always dropped out of the auction when the price exceeded their value, and market clearing prices were very near the Nash prediction.

3.2 Independent Private-Values Auctions

In an independent private-values auction, agents know the actual distribution generating the private values, that is, they are told x_0 as well as ε. Milgrom and Weber (1982a,b) demonstrated that if agents in a first-price affiliated-values auction are given information about the distribution of values, then the equilibrium price of the object will increase. Kagel *et al.* (1987) show that the Nash equilibrium bid function for first-price independent-values auctions is[10]

$$b(x_i) = \frac{n-1}{n} x_i + \frac{1}{n}(x_0 - \varepsilon). \tag{3}$$

Expected profits of the winning bidder can be calculated as[11]

$$E\pi = \frac{2\varepsilon}{n+1}. \tag{4}$$

Again, second-price auctions for independent-values have a dominant strategy of bidding one's true value. Hence, the equilibrium profits are the same as (2).

A striking result of auction theory, first noted by Vickrey (1961), is the revenue equivalence theorem. This states that when risk-neutral bidders have private values for an item, and when those values are completely independent, then the first-price and second-price auctions will yield the same expected revenue to the sellers. Therefore, (2) and (4) should be identical.

Kagel *et al.* (1987) tested the Milgrom and Weber prediction by providing subjects with x_0 after they bid in an affiliated-values auction, and thus transforming the auction into an independent-values auction. In theory this should cause bidders' profits to fall. They found that average profits across all treatments did fall, but that they did not fall in every treatment. In three of the five treatments average profits were virtually unchanged, even though bidders reacted strongly to the new information, and in one of the treatments profits actually increased. Only in one treatment did the

[10] In contrast to (1), (3) is the exact Nash bidding strategy. Hence, if the approximation given in (1) has any effect, we would expect that, relative to the independent values auctions, the approximation should hamper convergence to the predicted parameters in the common- and affiliated-values auctions. As we will see, this does not appear to be the case. Hence, any effect of the approximation, if it is present, would tend to work against our conclusions.

[11] Equilibrium expected profits can be found by substituting the expected value of the nth order statistic, x^n, into (3), and evaluating $x^n - b(x^n)$.

profits fall significantly. In a study of individual bidder behavior, Kagel and Levin (1985) reject all models they considered to explain the data, including asymmetric risk aversion.

3.3 Common-Values Auctions

In common-values auctions the bidders compete for an item that is of unknown value at the time of the auction, but of identical value to all bidders after the auction. Examples of such auctions include offshore oil leases and construction contracts. The auction is by sealed bid, and the highest bidder wins. The difficulty for bidders is that they must account for the fact that there is information in being the highest bidder. In particular, if bidders use similar monotonic bidding functions then the highest bidder is the one who most "over-estimated" the value of the item. A bidder who fails to account for this possibility may suffer a "winner's curse" since profits from winning may be unexpectedly low.

The structure of our simulated common-values auctions is the same as that used in the Kagel and Levin (1986) experiment. Let x_0 be the common value of the item. Each bidder, i, is given a signal of the item's value, x_i, where ε is the maximum possible estimation error. Hence, each bidder's signal is an unbiased estimate of the value of the item. Kagel and Levin (1986) solve for the Nash equilibrium bid function for this model, assuming that bidders know the process by which values and signals are generated. They find that the Nash equilibrium bid function is[12]

$$b(x_i) = x_i - \varepsilon. \tag{5}$$

The expected profits conditional on winning are[13]

$$E\pi = \frac{2\varepsilon}{n-1}.$$

Experiments by Kagel and Levin (1986) found that, in general, subjects suffer a winner's curse, especially in large groups (five to seven bidders). However, they find that small groups (four bidders) eventually learn to

[12] As before (see footnote 8), over the range $[1000 + \varepsilon, 2000 - \varepsilon]$, the exact bid function is $b(x_i) = x_i - \varepsilon + y(x_i)$. Again, our simplified bid function captures the relevant dimensions on this interval, and we therefore do not believe this simplification affects our results or conclusions. In fact, $y(x_i)$ is almost always extremely small relative to the Nash equilibrium bid.

[13] In Nash equilibrium the bidder with the highest signal will always be the highest bidder. Hence, the expected profits are the expected value of $b(x_i) - x_0$, conditioning on the fact that x_i is the highest signal.

bid profitably on average, although they still earn less than the Nash equilibrium prediction. The large groups continue to lose money, even with experience.[14] Dyer *et al.* (1989) found that these effects also exist among "professional" bidders.

4. Results of Auctions with AAA Bidders

We conducted all five auctions discussed above for both four- and eight-bidder groups. For each auction experiment we ran 20 trials, with each trial lasting 1000 generations. The algorithm searches over 2-parameter linear functions in each auction. For common- and affiliated-values auctions, the function's variables are x_i and ε, while for independent-values auctions the variables are x_i and $x_0 - \varepsilon$. The predicted Nash equilibrium values of the parameters on the bid functions and the expected profits (based on the mean value of ε of 250) are provided in the appropriate data tables. Note that for each auction there are two parameters that must be chosen by the AAA bidders. From this perspective, the parameter space over which the algorithm must search is identical for all auctions. This means that any difference in the ability of the algorithm to find Nash equilibria will not be attributable to the dimensions of the search space, but rather to the informational differences in the auctions.

Three different environments are explored for each experiment: *coevolution*, *full-feedback*, and *Nash-opponents*. In the coevolutionary environment strategies interact with other evolving strategies. Auctions are held in randomly formed groups, and payoffs equal the bidder's actual profit or loss.[15] This environment will be the focus of our study. We also explore two other environments to help evaluate our results. In the full-feedback environment, strings do not participate in auctions but instead each strategy receives a payoff equal to minus its squared deviation from the Nash equilibrium bid. This environment is used to test the algorithm's ability to solve the function fitting problem. The payoffs in the full-feedback condition are independent of the choices of other agents, and do not depend on winning auctions. In addition, the payoff space in the full-feedback condition is identical across all auctions; that is, accuracy is rewarded identically. In the Nash-opponents environment, each strategy

[14] These results have been replicated by Lind and Plott (1991).

[15] In this environment, two populations of 40 bidders are simultaneously evolved, where the two populations share no "genetic material." By having two populations coevolve against one another some endogeneity problems inherent in evolving only a single population are avoided. For instance, purely stochastic forces can cause a single population to rapidly converge towards a common string. This *genetic drift* effect was first noticed in natural populations.

bids in auctions consisting entirely of opponents who use the Nash equilibrium bidding strategy. If the strategy wins the auction, then its payoff equals its actual profit or loss, otherwise its payoff is zero. This differs from the coevolutionary condition in that each strategy evolves against a constant environment. Since the Nash equilibrium bid is the best response in an environment consisting entirely of other Nash equilibrium bidders, this condition tests the system's ability, in a stable environment, to find the best-response function.

4.1 Affiliated-Values Auctions

The results of the first- and second-price auctions with affiliated-values are listed in Table I. The numbers in this table, and in all subsequent tables, list the average values of the last 25 generations in each trial, averaged over all 20 trials. The standard deviations refer to the differences across trials.[16] We will refer to the parameters of the bid function as a pair (β_1, β_2) where β_1 refers to the parameter on x_i in the bidding function.

We find that the model tends to support theoretical predictions in first-price auctions under coevolution. For four-bidder groups the predicted values of the bid function are $(1, -0.5)$, while the model finds parameters of $(0.98982, -0.46417)$. The standard errors on these parameters are small, indicating wide agreement on the bidding function across trials. Moreover, average profits are 113.42, which is 90% of the Nash prediction of 125. There is similar success with the eight-bidder groups. The predicted parameters are $(1, -0.25)$, while the GA finds $(0.99540, -0.24694)$. Again, the standard errors across trials are small. However, eight-bidder groups earn average profits of 36.79, which is only 58% of the Nash prediction.

We see that the algorithm has largely been able to converge to the Nash equilibrium. The small errors that remain are only slightly costly to the four-bidder groups, and more costly to the larger groups. This leads us to ask whether these remaining errors can be attributed to the algorithm per se, or whether they can be attributed to the kind of informational environment present in a coevolutionary game. We can examine this question with the Nash-opponents and full-feedback conditions. We see

[16] Note that the only meaningful comparison here is *across* rather than within trials. The reason is that the standard deviations within trials are very small by the 1000th generation, usually on the order of 0.0001 to 0.05. This is due to the fact that the probability of mutation has decayed from its initial 8% to about ½% by the last generation, coupled with the "founding effect" phenomenon inherent in small populations. Thus, the members of a given population are likely to become very homogeneous over time (see Goldberg (1989) for further details about convergence in genetic algorithms). However, if different populations converge to very similar strings, then this implies the existence of a very large and strong basin of attraction for the learning dynamics. Hence, variance across population trials will be most indicative of the algorithm's ability to find a Nash equilibrium.

TABLE I

RESULTS FOR AFFILIATED-VALUES AUCTIONS

Information	Group size	Bidding function		Average profit[a]
		x_i	ε	
Predicted				
First-price	4	1.000	−0.500	125.0
	8	1.000	−0.250	62.5
Second-price	4	1.000	0.000	100.0
	8	1.000	0.000	55.5
Coevolving				
First-price	4	0.98982	−0.46417	113.42
		(0.00725)	(0.06139)	(8.77)
	8	0.99540	−0.24694	36.79
		(0.00347)	(0.02816)	(10.99)
Second-price	4	0.99955	0.00008	96.66
		(0.01034)	(0.07861)	(7.99)
	8	0.99754	0.04897	48.16
		(0.00475)	(0.04563)	(4.39)
Nash-opponents				
First-price	4	0.99116	−0.46892	122.85
		(0.00755)	(0.05013)	(13.37)
	8	0.99303	−0.21685	42.07
		(0.00395)	(0.01925)	(19.02)
Second-price	4	0.99821	0.00881	94.92
		(0.00860)	(0.06282)	(11.89)
	8	0.99471	0.03198	53.23
		(0.00471)	(0.04163)	(8.06)
Full-feedback				
First-price	4	0.99842	−0.49173	125.54
		(0.00776)	(0.03447)	(4.55)
	8	1.00130	−0.25634	61.88
		(0.00515)	(0.02009)	(4.45)
Second-price	4	1.00107	−0.00495	100.11
		(0.00290)	(0.01620)	(0.29)
	8	1.00004	−0.00391	55.50
		(0.00280)	(0.01648)	(0.15)

Note. Mean values and standard deviations (in parentheses) are over the last 25 generations.

[a] Profits reported for the Nash auctions are conditional on winning the auction. Profits for the full-feedback auctions are calculated for each simulation as expected values, using the average strategy over the last 25 generations. Means and standard deviations are calculated across simulations.

that in the Nash-opponents condition the model performs about as well at finding the bid function, but does better as measured in profits. This is true for both four- and eight-bidder auctions. This difference can probably be ascribed to the fact that the variance of the bid functions across auctions

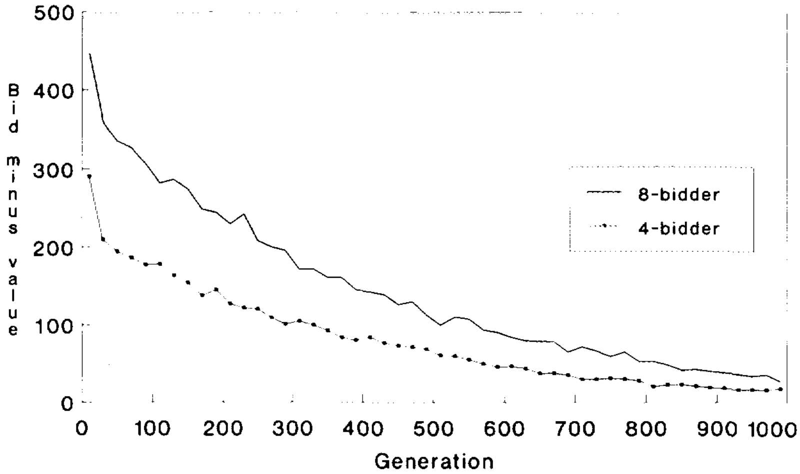

FIG. 1. Second-price affiliated-values auctions. Average of every 20 generations.

is smaller. Turning to the full-feedback condition, we see that performance improves even more. In both four- and eight-bidder auctions, the bid functions are almost exactly those predicted by the Nash equilibrium.

These results indicate that, while the coevolutionary environment does not generate the exact Nash prediction, it is nonetheless quite successful at finding the Nash bidding strategies in affiliated-values auctions. This is consistent with the experimental finding that Nash equilibrium bidding organizes the data better than other models, although errors still remain. The fact that the full-feedback auction can make the small improvements necessary to find the exact Nash bidding strategy indicates that the failure of the model to do even better in the coevolutionary environment cannot be blamed on the algorithm per se—in ideal feedback conditions, the algorithm can find the optimum. Instead, these results indicate that the small errors of the coevolutionary condition appear to be due to the fact that the feedback is less than ideal.

Turning to the second-price auctions we find similar results. For both four- and eight-bidder groups the coevolutionary strategies are extremely close to the Nash prediction of (1, 0). Moreover, average profits are 96% of the Nash prediction for the four-bidder group, and 87% of the Nash prediction for the eight-bidder group. The Nash-opponents condition performs about as well as the coevolution condition, with slightly less variance on the bid functions, while the full-feedback condition performs the best by attaining almost the exact Nash strategy.

Figure 1 illustrates the evolution of bids in the four- and eight-bidder second-price coevolving auctions. The figure shows the average deviation of the winning bid from the dominant strategy bid over all 1000 generations (averaged over every 20 generations). The deviations converge to the dominant strategy from above. We can explain this by examining the dynamics of adaptive learning. A strategy that bids beneath the dominant strategy will increase both its absolute and relative performance by raising its bid to the dominant strategy level. However, if a bidder is bidding above the dominant strategy, then moving to the dominant strategy will increase the bidder's absolute performance, but may reduce its relative performance. This is because the other strategies will now earn higher profits when this bidder is the second highest bidder. Hence, many situations exist in which a bidder will reduce its chance for survival by lowering its bid to the dominant strategy level. That is, from an evolutionary standpoint, truthful bidding may not always be dominant. Thus evolutionary forces may quickly eliminate under-bidders (who win no auctions) as well as the extreme over-bidders (who make losses), but may be slow to eliminate moderate over-bidders. This implies that evolutionary learning is likely to lead to convergence to the dominant strategy from above. Again, this is consistent with the experimental results.

4.2. *Independent Private-Values Auctions*

The results of the independent-values auctions are given in Table II. In the four-bidder coevolution auctions the average parameters are (0.75817, 0.21890), as compared to the predicted values of (0.75, 0.25). Furthermore, the profits are 81% of the predicted value. However, the standard errors on the bid parameters and profits are relatively large. Hence, even though the bid functions meet the prediction on average, there are very few trials that are close to the Nash equilibrium functions. Turning to the eight-bidder auctions, the performance of the coevolution condition erodes further. Again, the standard deviations are relatively large, but the average values of the bid function are not at all near the Nash prediction of (0.875, 0.125). Moreover, profits are almost zero. In fact, in 8 of 20 trials the average profits are actually negative. This means that in the eight-bidder auctions, the model is often not capable of finding functions that, on average, bid less than their value. In the second-price coevolution auctions, the parameter values are much farther from the predicted values than in the affiliated-values auctions, especially the second parameter. Moreover, the standard deviations are again very large. This is true for both four- and eight-bidder groups. Finally, there is a clear divergence from the revenue equivalence theorem, with bidders making more profits in the second-price auctions, especially with eight bidders.[17]

[17] We also conducted auctions that allowed for the choice of three parameters, that is bid functions of $b_i = \alpha_1 x_i + \alpha_2 x_0 + \alpha_3 \varepsilon$, and obtained similar results. In fact, the GA performed less well in this more general environment. For sake of brevity, and to maintain consistency in the search spaces across auctions, we only report the results of two-parameter bid functions.

TABLE II

RESULTS FOR INDEPENDENT-VALUES AUCTIONS

Information	Group size	Bidding function		Average profit[a]
		x_i	$x_0 - \varepsilon$	
Predicted				
First-price	4	0.750	0.250	100.0
	8	0.875	0.125	55.5
Second-price	4	1.000	0.000	100.0
	8	1.000	0.000	55.5
Coevolving				
First-price	4	0.75817	0.21890	81.06
		(0.19865)	(0.24443)	(20.42)
	8	0.93919	0.02677	0.94
		(0.12925)	(0.17177)	(18.95)
Second-price	4	1.11302	−0.14232	90.10
		(0.20634)	(0.25546)	(8.40)
	8	1.08450	−0.11753	41.29
		(0.11450)	(0.16327)	(5.75)
Nash-opponents				
First-price	4	0.94336	−0.03517	86.70
		(0.11131)	(0.17945)	(16.03)
	8	1.00255	−0.07987	30.21
		(0.05232)	(0.08961)	(23.96)
Second-price	4	1.13557	−0.18689	100.10
		(0.17900)	(0.23904)	(12.72)
	8	1.14978	0.22652	46.52
		(0.16246)	(0.25641)	(7.81)
Full-feedback				
First-price	4	0.82791	0.15733	87.29
		(0.23577	(0.28406)	(33.73)
	8	0.82793	0.18357	62.03
		(0.22899)	(0.27461)	(43.99)
Second-price	4	1.04453	−0.05090	104.45
		(0.29955)	(0.36125)	(29.96)
	8	0.99099	0.01149	55.00
		(0.26940	(0.32228)	(14.95)

Note. Mean values and standard deviations (in parentheses) are over the last 25 generations.

[a] Profits reported for the Nash auctions are conditional on winning the auction. Profits for the full-feedback auctions are calculated for each simulation as expected values, using the average strategy over the last 25 generations. Means and standard deviations are calculated across simulations.

The model's difficulty in finding optimal bidding rules in independent-values auctions is in stark contrast to its relative success in the previous auctions. We can examine this more carefully by looking at the Nash-opponents and full-feedback conditions. As can be seen in Table II, neither the Nash-opponents or full-feedback auctions are significantly closer to the Nash equilibrium than are the coevolution auctions. This is true for both first- and second-price auctions and for both large and small groups. Of particular note are the relatively large standard errors on the parameters of the bid functions for all three information conditions.

One hypothesis for why subjects do not reach Nash equilibria in auction experiments is that the payoff space in independent-values auctions may be very flat, hence there may be little incentive for improved bidding (Harrison, 1989). There is some evidence for that here. For the four-bidder coevolving auctions, only 10 of the 40 bid functions are within 0.05 of the Nash strategies on both parameters. However, when evaluating the average bid in each trial, we find that 38 of 40 bid functions make bids that are, on average, within 5% of the Nash bid, even though the functions may be very different.[18] In the four-bidder second-price auctions, only 8 of the 40 bid functions are close to Nash, but 39 of them bid within 3% of the Nash bid on average.[19] Similar patterns emerge in the eight-bidder auctions. Hence, there may be a large family of strategies that, over many private values, do about as well as Nash. Significant movements toward Nash behavior in function space may translate to only small improvements in payoff space. This may make the convergence to the Nash equilibrium particularly slow.

To examine the conjecture that flatness of the payoff space can explain the failure to learn equilibrium play, we assumed that the population is risk averse rather than risk neutral, and hence dislikes the variance that is inherent in the non-Nash bidding strategies. We reran the first- and second-price coevolutionary auctions using a risk-averse utility function for the bidders. We assume, as in Kagel *et al.* (1987), a constant relative risk-averse utility function $U = \pi^r$, where r is the risk-aversion parameter. On the basis of their experiments, we chose $r = 0.5$. To account for the possibility that profits in the actual auctions may sometimes be negative, we linearized the utility function in the neighborhood of zero. Hence, the AAA bidders had the utility function

$$U_i = \begin{cases} \pi^{0.5} & \text{if } \pi \geq 0.01; \\ 5\pi + 0.05 & \text{if } \pi < 0.01. \end{cases}$$

[18] For instance, a bid function that is "close," with parameters (0.74427, 0.24825), bids on average 1426.72, while another bid function that is not close, with parameters (0.48912, 0.55425), bids on average 1426.49.

[19] For example, a close function with parameters (1.0268, −0.03770) makes an average bid of 1493.14, while a bid function with parameters (1.2611, −0.3184) makes an average bid of 1493.64.

TABLE III

INDEPENDENT-VALUES AUCTIONS WITH RISK AVERSE BIDDERS

Auction	Group size	Bidding function x_i	Bidding function $x_0 - \varepsilon$	Average profit
Predicted				
First-price	4	0.85714	0.14286	57.14
	8	0.93333	0.06667	26.14
Second-price	4	1.000	0.00000	100.00
	8	1.000	0.00000	55.50
Coevolving				
First-price	4	0.93108	0.4027	24.40
		(0.10451)	(0.14704)	(23.22)
	8	0.99476	−0.02763	27.43
		(0.05532)	(0.11483)	(21.54)
Second-price	4	1.03363	−0.06133	113.70
		(0.06557)	(0.09909)	(14.63)
	8	1.00832	−0.02394	55.67
		(0.04950)	(0.07298)	(10.34)

Note. Mean values and standard deviations (in parentheses) are over the last 25 generations.

Kagel *et al.* (1987) show that in first-price auctions the Nash equilibrium bidding function is $b_i = [(n - 1)/(n - r)]x_i + [(1 - r)/(n - r)](x_0 - \varepsilon)$. For four-bidder groups the predicted parameters are (0.8572, 0.1428), with profits of 57.14. For eight-bidder groups the predicted parameters are (0.9333, 0.0667), with profits of 26.14. For second-price auctions, bidding one's value is still dominant, so the predictions are identical to the risk-neutral case.

The results of the auctions with risk aversion are listed in Table III. The addition of risk aversion brought only slight improvements in the first-price auctions. While the standard deviations are reduced somewhat, the bid functions are still fairly inaccurate. However, for the second-price auctions the average bid functions become much more accurate, while the standard erorrs become relatively small. Hence, risk aversion appears to help bidders in second-price auctions. Note also that the prediction that risk aversion should reduce bidder profit was upheld in the four-bidder auctions, but contradicted in the eight-bidder auctions. This matches the contradictory nature of the evidence on risk aversion with human experiments.

This suggests that the flatness of the payoff space alone cannot explain the results in the first-price auctions. This view is also supported by the results of the full-feedback condition. In this condition the payoff space is not flat, but is relatively steep. Moreover, the shape of the payoff space is identical to that of the affiliated-values auction, yet in that auction the algorithm has little difficulty. The difference between the two is the function space. With affiliated-values the algorithm fits parameters to (x_i, ε), while with independent-values it fits parameters to $(x_i, x_0 - \varepsilon)$. One important difference is that the correlation between x_i and ε is $\rho(x_i, \varepsilon) = 0$, while the correlation between x_i and $x_0 - \varepsilon$ is $\rho(x_i, x_0 - \varepsilon) = \rho(x_i, x_0) = 0.77$ (this correlation has often been higher in human experiments). Hence, one possibility is that it is difficult for adaptive agents to separate the effects of x_i and $x_0 - \varepsilon$. As a result, a great deal of experimentation is necessary to learn the appropriate bid function.

The last conjecture can be examined by writing the objective function in a way that avoids the correlation. Recall that a general application of the theory actually suggests a bid function with three parameters, $b_i = \beta_1 x_1 + \beta_2 x_0 + \beta_3 \varepsilon$. Our initial simulations can be thought of as artificially restricting the parameters to be $\beta_2 = -\beta_3$, which is the Nash equilibrium solution. This maintains the dimension and quality of the search space relative to the affiliated values auction. Instead, let us impose the restriction $\beta_2/\beta_1 = 1/(n - 1)$. Hence, we can rewrite the bid function $b_i = \beta_1 x_i + \beta_2(x_0 - \varepsilon)$ as $b_i = \beta_1(x_i + \alpha x_0) - \beta_2 \varepsilon$, where $\alpha = \beta_2/\beta_1$. Now $\rho(x_i + \alpha x_0, \varepsilon) = 0$. Looking at first price auctions, we set $\alpha = \frac{1}{3}$ for the four-bidder groups, and $\alpha = \frac{1}{7}$ for the eight-bidder groups. This bid function does not change the payoff space for the independent-values auctions, but reorganizes the function space to eliminate the correlation.[20] We find that for four-bidder full-feedback groups the model converges to the optimal values ($\beta_1 = 0.752$, SD = 0.005; $\beta_2 = 0.257$, SD = 0.034). For eight-bidder groups, sixteen of the twenty trials found the optimum ($\beta_1 = 0.879$, SD = 0.007; $\beta_2 = 0.151$, SD = 0.051), however, four trials strayed to other values. This tends to confirm the hypothesis that the highly correlated information in independent-values auctions, rather than the flat payoff space, is complicating adaptive learning.

4.3. Common-Values Auctions

The results for common-values auctions are listed in Table IV. Looking first at the four-bidder coevolutionary auctions, we see that the bid functions are in the neighborhood of the Nash equilibrium strategies, with

[20] Reorganizing the bid function in other ways does not eliminate the correlation. For instance, rewriting the function as $b_i = \beta_1(x_0 - \varepsilon) + \beta_2(x_i - (x_0 - \varepsilon))$, yields a correlation of $\rho(x_0 - \varepsilon, x_i - (x_0 - \varepsilon)) = 0.23$. Even though this yields a smaller correlation, we prefer our original formulation (3) since it makes the fewest assumptions on the cognition of our adaptive agents, and, moreover, it more accurately represents the way the information is actually presented in experiments.

TABLE IV

RESULTS FOR COMMON-VALUES AUCTIONS

Information	Group size	Bidding function x_i	Bidding function ε	Average profit[a]
Predicted	4	1.000	−1.000	100.0
	8	1.000	−1.000	55.5
Coevolving	4	0.98238	−0.87873	81.90
	8	(0.01993)	(0.12753)	(8.13)
		0.99360	−0.94889	18.20
		(0.00547)	(0.04137)	(15.73)
Nash-Opponents	4	0.86627	−0.34417	112.71
		(0.10645)	(0.48812)	(18.80)
	8	0.95150	−0.74271	35.10
		(0.06015)	(0.28120)	(17.80)
Full-Feedback	4	0.99876	−0.98581	98.51
		(0.00452)	(0.02472)	(3.42)
	8	1.00040	−0.99486	54.05
		(0.00226)	(0.01196)	(2.91)

Note. Mean values and standard deviations are over the last 25 generations.

[a] Profits reported for the Nash auctions are conditional on winning the auction. Profits for the full-feedback auctions are calculated for each simulation as expected values, using the average strategy over the last 25 generations. Means and standard deviations are calculated across simulations.

small standard errors, but are much farther from the prediced levels than are the affiliated-values auctions. In terms of profits, the coevolutionary strategies earn 82% of the Nash prediction. In all 20 trials the AAA bidders earned positive profits, with the maximum at 96.1 and the minimum at 72.2. Turning to the eight-bidder coevolutionary auctions we see that the parameters of the bid function are closer to the Nash equilibrium values than in the four-bidder auctions. Again, the standard deviations of these parameters are small. Average profits are 18.20, which is only 33% of the Nash equilibrium amount, and in two of the twenty trials the AAA bidders actually made losses. Overall profits varied from 51.47 to −10.59. Hence, even though the bidders in eight-bidder auctions appear to make smaller errors, these errors are more costly.

Kagel and Levin (1986), and Dyer *et al.* (1989) showed that, with sufficient experience, bidders in four-bidder groups are able to learn to bid profitably, even though they make smaller profits than predicted, while larger groups are generally unable to avoid losses. We find a similar difference in profitability in the coevolutionary auctions. Kagel and Levin (1986) conjectured that learning is "situationally specific," and that for large groups the situation leads to more aggressive bidding, contrary to the Nash equilibrium prediction. This indicates that perhaps learning is more difficult in larger groups. Our results are consistent with Kagel and Levin (1986) in that bidders in the larger groups actually do learn to bid more aggressively. This can be seen in the much lower profits earned by the eight-bidder groups. It is interesting to note, however, that our results indicate that the eight-bidder environment does not necessarily make learning more difficult. In fact, the algorithm is more accurate at finding the Nash bid function in eight-bidder auctions than in four-bidder auctions. Instead, the main difference appears to be that small errors are more costly in larger groups, so that bidders have to be more accurate to achieve the same levels of profitability.[21] These results indicate that the patterns observed in experiments are consistent with evolutionary learning, but they do not support the conjecture that larger groups make learning more difficult.

Looking at the Nash-opponents and full-feedback conditions we can see how learning in common-values auctions may differ from affiliated-values auctions. In the full-feedback auction for both four- and eight-bidder groups, we see that, as above, the model is able to find the Nash function with good precision. This should be expected, since the function space and the payoff space are identical to affiliated values auctions. Turning to the Nash-opponents auctions, however, we see that the AAA bidders find parameters that are farther from the Nash strategies than are the coevolutionary strategies, and the variances on the parameters are relatively large. The bidders in the Nash-opponents condition actually make more profits when they win auctions than do the coevolutionary strategies, and they make more than predicted by equilibrium Nash behavior. Examining the individual bids, we found that in the Nash-opponents condition the algorithm chooses functions that systematically bid beneath the Nash equilibrium bid. As a result, the AAA bidders win less often, but earn much more when they do, while the Nash bidders win more often, and earn slightly less. The net effect is that the AAA bidders make almost the same profits per auction as the Nash bidders.[22] This indicates that perhaps adaptively learning the best response is difficult in common-

[21] For instance, in the four-bidder auctions, bids exceed the Nash equilibrium prediction by 11.65 on average, but earnings are 80% of Nash equilibrium profits. The bid function for the eight-bidder auction overbid by only 3.15 on average, yet earn only 33% of the Nash equilibrium profits. (These values were calculated at the mean values of ε and x_0.)

[22] For example, we simulated 1000 auctions in which the AAA strategy bid against three other Nash strategies. While the AAA strategy won only 158 auctions, it earned an average of 142.40 when it won the auction. The Nash strategies earned an averaged of 84.80 when they won. However, the earnings per auction were 22.5 for the AAA strategy and 23.8 for the Nash.

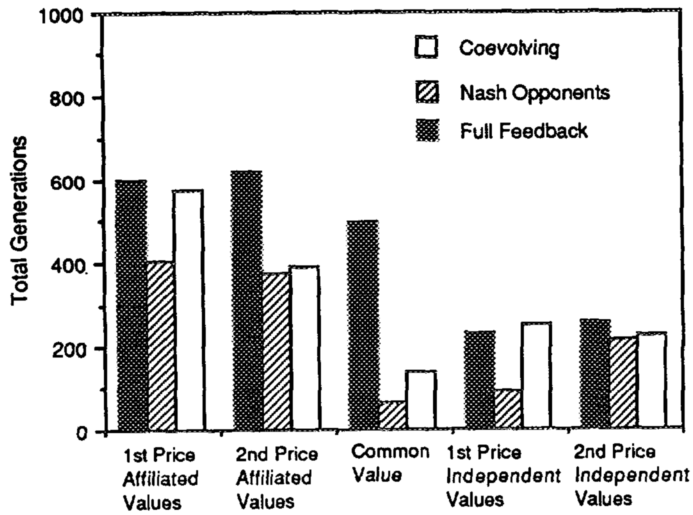

FIG. 2. Number of generations close to Nash strategy for four-bidder groups.

values auctions because there may be several bid functions that do about as well as Nash on average, even though the variance of the returns is very high. As a result, it may take many generations for the algorithm to discover the innovations in the bid function that would take advantage of the small gains available. Another possible explanation arises from the fact that the best strategy in the population may sometimes make losses. This means that in the short run the "best" strategy will sometimes do worse than others, even though in the long run it will always do better. Hence, chance losses may bias the search mechanism away from Nash strategies.[23]

5. Relative Difficulty of the Auction Environments

One measure of the difficulty of learning is how frequently the model can develop bid functions that are "close" to the Nash function. To examine the relative difficulty of learning, define a strategy as "close" if *both* parameters are within a range of 0.075 to the Nash parameters (comparable results obtain for other ranges). Figures 2 and 3 show the

[23] It does not, however, appear that the outcome in the Nash-opponents condition can be attributed to the simplification in the Nash strategy. In particular, the approximation does not hamper the search process in the coevolutionary auctions, leading us to suspect that it is the feedback, rather than the approximation, that is responsible for this result.

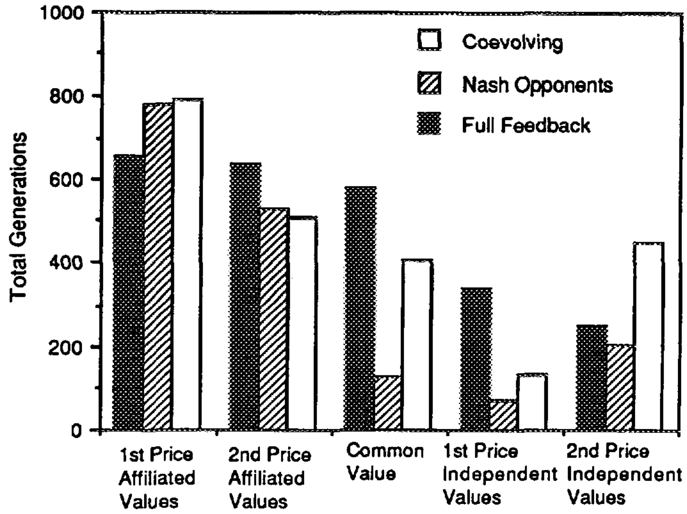

FIG. 3. Number of generations close to Nash strategy for eight-bidder groups.

average number of generations in which the mean evolving strategy is close to the Nash strategy.

Under full-feedback the algorithm is able to find parameters close to the Nash equilibrium values in the affiliated and common-value auctions in 50 to 60% of the generations, while in the two independent-values auctions the evolving parameters are close to Nash strategies in only 25 to 33% of the generations. This provides further evidence that searching the parameter space is much harder in the independent-values auctions than in the other two.

In the Nash-opponents and coevolutionary environments we see that the probability of convergence is greatest in the affiliated-values auctions. The difficulty of convergence in the independent-values auctions is not surprising since convergence is difficult even under full-feedback. The relative lack of convergence in the common-value auction cannot be so easily explained since under full-feedback, the convergence in the common-value auction is only slightly less than that observed in the affiliated-values auctions. Instead, as conjectured in the previous section, it appears that the adverse feedback may be having an effect on the adaptive dynamics.

Notice that in nine of the ten auctions the strategies in the coevolutionary environment converge more often than those in the Nash-opponents environment. One might expect that adaptation to a best response against Nash opponents should be easier than adaptation against non-Nash opponents. However, an environment consisting of all Nash opponents will very likely result in zero or negative payoffs to non-Nash players, whereas agents in a coevolving environment will have more opportunity for positive payoffs. Thus, in the coevolving environment the system gets a lot of information about areas of the search space that are likely to be productive, while in the Nash-opponents condition the only information forthcoming is about nonproductive areas. Given the search task and the scarcity of productive parts of the strategy space, coevolution may facilitate learning Nash equilibria better than the stationary environment of Nash equilibrium players. This appears to coincide well with the common experience that one can learn best in environments that are only slightly challenging—a beginner will not learn to play tennis by initially facing John McEnroe.

6. Conclusions

Experiments on auctions find that subjects make systematic bidding errors that cannot be explained within the context of Nash or risk-averse Nash bidding models. This paper has considered these errors using a model of adaptive learning based on a genetic algorithm. We find that artificial adaptive agents exhibit many of the same bidding patterns observed in auctions with humans. In particular, while subjects in second-price auctions do better than those in first-price auctions, bidders tend to over-bid in all auctions. Independent values appear to lead to higher variances in bidding relative to affiliated values. It appears that the additional information in independent-values auctions serves to complicate the bidding task. This is even true in second-price auctions. Adding risk aversion to the bidders improves the performance in second-price independent-values auctions, but provides contradictory results in first-price independent-values auctions. Finally, in common-values auctions bidders in small groups do a better job of avoiding the winner's curse, just as in human experiments. We find that adaptive learning can provide useful insights into bidder behavior and can reconcile existing theoretical and experimental results.

More specifically, we conclude the following:

1. The observed difference between small and large groups in common-values auctions may have more to do with how the probability of a winner's curse is increased in large groups than with the ability of subjects to learn Nash bidding. Our results indicate that the size of the group alone does not affect the ability of the bidders to learn Nash equilibrium bidding. In fact, the eight-bidder groups actually outperformed the four-bidder groups in strategy space. However, like human experiments, the larger groups suffered a larger winner's curse, indicating that the observed differences may be attributable to the fact that a given error is more costly in

larger groups, rather than to differences in learning or rationality across groups.

2. Learning the dominant strategy in second-price auctions is not simple, and adaptive learning can lead to systematic overbidding. Because increasing one's bid in a second-price auction creates a negative externality on other bidders, increasing one's bid can sometimes increase one's relative performance, even if that bid exceeds the dominant strategy. This effect may tend to increase the fitness of over-bidding strategies relative to dominant strategy bidders or under-bidders. As a result, dominant strategy bidding will tend to converge from above.

3. Bidding errors in independent-values auctions cannot be attributed only to the flatness of the payoff space. Comparing the full-feedback condition of the independent-values and affiliated-values auctions, we see that the pure function fitting problem is much easier in affiliated-values auctions, even though the shape of the payoff space is identical for both auctions. This implies that there must be something in the function space that differs between the two. Because information given in independent-values auctions is highly correlated it is difficult for bidders to distinguish the effect of weight placed on x_i from that placed on $x_0 - \varepsilon$. Hence, evolving toward Nash equilibrium may be more difficult in independent-values auctions.

4. Affiliated-values auctions facilitate Nash equilibrium bidding better than either independent- or common-values auctions. Compared to independent-values auctions, correlation between the pieces of information given the bidder in the affiliated-values auctions is much lower. Hence, it is easier for bidders to distinguish the effects of different parameters on the bid function. Compared to the common-values auctions, the information is of the same quality, but the feedback is superior. In common-values auctions the "best" strategy in the population can sometimes make losses. Chance events can make a strategy that is better *ex ante* look worse *ex post*. Private-values auctions do not suffer from this adverse and contradictory feedback, again making it easier to evolve toward Nash equilibrium bidding rules.

5. Buyer's profits are lower—and hence seller's profits are higher—in first-price independent-values auctions than in second-price auctions. This contradicts the revenue equivalence theorem. If this result generalizes to real world auctions, then this provides one more hypothesis for why second-price auctions are less common in practice than sealed bid first-price auctions (see Rothkopf *et al.* 1990); bidding errors in the favor of sellers may lead them to choose first-price auctions.

6. Auctions are very difficult environments for evolutionary learning. Even in the easiest auctions, the AAA bidders are only able to come "near" the Nash solution in about 600 of the 1000 generations on average,

while for the hardest auctions this number is as low as 134. Compare this to Miller and Andreoni's (1991) examination of public goods experiments in which adaptive agents converged to the Nash Equilibrium in under 40 generations. The relative difficulty of auctions appears to be attributable to both the poor quality and low amount of feedback inherent in these environments. Adaptive systems must spend many generations eliminating bad strategies before they can begin to significantly strengthen the good strategies. Furthermore, when a bidder wins profitably, it may be due to a savvy strategy or simply to chance. Hence, a great deal of experimentation and innovation is necessary to determine what, if any, improvements are needed. This again slows the search for better strategies. Most human experiments give subjects experience in about 18 to 30 auctions. While this amount of experience is quite sufficient by most experimental standards, our results indicate that, in general, auctions are a relatively difficult environment for adaptive learning. Thus, experiments designed to test the ability of bidders to learn optimized bidding rules may require much longer exposures of subjects to the auction games.

7. Participating in an environment of other evolving strategies appears to facilitate learning a Nash equilibrium better than trying to learn a best response to a fixed environment of Nash equilibrium players. There are several reasons for this. An environment of Nash equilibrium players is very challenging for adaptive learning. The evolving strategies are more likely to receive zero or negative payoffs than positive payoffs, and thus the task of searching for sets of good strategies is enormously complicated relative to coevolving systems.

We have shown that adaptive learning may provide an explanation for the divergence between theoretical and experimental results in auction markets. By performing computational experiments on generalized adaptive algorithms, we can begin to characterize dynamic learning behavior in these auctions. Along with reconciling existing theory and experiments, this approach also suggests some new research directions. Central to our work is the ability to study a generic adaptive system. Computational methods allow the feasible analysis of simple systems of interacting adaptive agents that lead to complex stochastic dynamics. The addition of adaptive benchmarks to the large amount of existing theoretical and experimental results enhances current research efforts, and has broad potential throughout economic analysis.

References

AXELROD, R. (1987). "The Evolution of Strategies in the Iterated Prisoner's Dilemma," in *Genetic Algorithms and Simulated Annealing* (L. Davis, Ed.). Los Altos, CA: Morgan Kaufmann.

BINMORE, K., AND SAMUELSON, L. (1992). "Evolutionary Stability in Repeated Games Played by Finite Automata," *J. Econ. Theory* **57**, 278–305.

BOYLAN, R. (1990), "Evolutionary Equilibria with Mutation," California Institute of Technology working paper.

COX, J., ROBERSON, B., AND SMITH, V. (1982). "Theory and Behavior of Single-Object Auctions," *in Research in Experimental Economics* (V. L. Smith, Ed.). Greenwich: JAI Press.

COX, J., SMITH, V., AND WALKER, J. (1983). "Tests of Heterogeneous Bidder's Theory in First Price Auctions," *Econ. Lett.* **12**, 207–212.

COX, J., SMITH, V., AND WALKER, J. (1985). "Experimental Development of Sealed-Bid Auction Theory: Calibrating Controls for Risk Aversion," *Amer. Econ. Rev. Papers Proc.* **75**, 160–165.

COX, J., SMITH, V., AND WALKER, J. (1992). "Theory and Misbehavior in First-Price Auctions: Comment," *Amer. Econ. Rev.* **82**, 1392–1412.

CRAWFORD, V. P. (1991). "An 'Evolutionary' Explanation of Van Huyck, Battalio, and Beil's Experimental Results on Coordination," *Games Econ. Behav.* **3**, 25–59.

DYER, D., KAGEL, J. H., AND LEVIN, P. (1989). "A Comparison of Naive and Experienced Bidders in Common Value Offer Auctions: A Laboratory Analysis," *Econ. J.* **99**, 108–115.

FRIEDMAN, D. (1992). "Theory and Misbehavior in First-Place Auctions: Comment," *Amer. Econ. Rev.* **82**, 1374–1378.

GOLDBERG, D. (1989). *Genetic Algorithms in Search, Optimization, and Machine Learning*, New York: Addison-Wesley.

HARRISON, G. W. (1989). "Theory and Misbehavior in First-Price Auctions," *Amer. Econ. Rev.* **79**, 749–763.

HOLLAND, J. (1975). *Adaptation in Natural and Artificial Systems*, Ann Arbor: University of Michigan Press.

HOLLAND, J. H., AND MILLER, J. H. (1991). "Artificial Adaptive Agents in Economic Theory," *Amer. Econ. Rev. Papers Proc.* **81**, 365–370.

KAGEL, J. H., AND DYER, D. (1988). "Learning in Common Value Auctions," in *Experimental Games and Markets* (R. Tietz, W. Albers, and R. Selton, Eds.). Berlin: Springer-Verlag.

KAGEL, J. H., HARSTAD, R. M., AND LEVIN, D. (1987). "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study," *Econometrica* **55**, 1275–1304.

KAGEL, J. H., AND LEVIN, D. (1985). "Individual Bidder Behavior in First Price Private Values Auctions," *Econ. Lett.* **19**, 125–128.

KAGEL, J. H., AND LEVIN, D. (1986). "The Winner's Curse and Public Information in Common Value Auctions," *Amer. Econ. Rev.* **76**, 894–920.

KAGEL, J. H., LEVIN, D., AND HARSTAD, R. M. (1988). "Judgment, Evaluation and Information Processing in Second-Price Common Value Auctions," University of Pittsburgh working paper.

KAGEL, J. H., AND ROTH, A. (1992). "Theory and Misbehavior in First-Price Auctions: Comment," *Amer. Econ. Rev.* **82**, 1379–1391.

KOLLMAN, K., MILLER, J. H., AND PAGE, S. E. (1992). "Adaptive Parties in Spatial Elections," *Amer. Polit. Sci. Rev.* **86**, 929–937.

LIND, B., AND PLOTT, C. R. (1991). "The Winner's Curse: Experiments with Buyers and Sellers," *Amer. Econ. Rev.* **81**, 335–346.

MARIMON, R., MCGRATTAN, E., AND SARGENT, T. J. (1990). "Money as a Medium of Exchange in an Economy with Artificially Intelligent Agents," *J. Econ. Dynam. Control* **14,** 329–373.

MERLO, A., AND SCHOTTER, A. (1992). "Theory and Misbehavior in First-Price Auctions: Comment," *Amer. Econ. Rev.* **82,** 1413–1425.

MILGROM, P. R., AND WEBER, R. J. (1982a). "A Theory of Auctions and Competitive Bidding," *Econometrica* **50**(5), 1089–1122.

MILGROM, P. R., AND WEBER, R. J.. (1982b). "The Value of Information in a Sealed Bid Auction," *J. Math. Econ.* **10**(1), 105–114.

MILLER, J. H. (1986). "A Genetic Model of Adaptive Economic Behavior," University of Michigan working paper.

MILLER, J. H. (1995). "The Coevolution of Automata in the Repeated Prisoner's Dilemma," *J. Econ. Behav. Organ.* to appear.

MILLER, J. H., AND ANDREONI, J. (1991). "Can Evolutionary Dynamics Explain Free Riding in Experiments?," *Econ. Lett.* **36,** 9–15.

ROTHKOPF, M. H., TEISBERG, T. J., AND KAHN, E. P. (1990). "Why are Vickrey Auctions so Rare?," *J. Polit. Econ.* **98,** 94–109.

RUST, J., MILLER, J. H., AND PALMER, R. (1992). "Behavior of Trading Automata in a Computerized Double Auction Market," in *The Double Auction Market: Institutions, Theories, and Evidence* (D. Friedman and J. Rust, Eds.), pp. 155–198. Reading, MA. Addison Wesley.

RUST, J., MILLER, J. H., AND PALMER, R. (1994). "Characterizing Effective Trading Strategies: Insights from a Computerized Double Auction Tournament," *J. Econ. Dynam. Control* **18,** 61–96.

SCHAFFER, J. D., CARUANA, R., ESHELMAN, L., AND DAS, R. (1989). "A Study of Control Parameters Affecting Online Performance of Genetic Algorithms for Function Optimization," in *Proceedings of the Third International Conference on Genetic Algorithms* (J. D. Schaffer, Ed.). San Mateo, CA: Morgan Kaufmann.

VICKREY, W. (1961). "Counterspeculation, Auctions, and Competitive Sealed Tenders," *J. Finance* **16**(1), 8–37.