Chapter 17

INDIVIDUAL RATIONALITY AS A PARTIAL IMPEDIMENT TO MARKET EFFICIENCY

Allocative Efficiency of Markets with Smart Traders

Shu-Heng Chen
AI-ECON Research Center, Department of Economics, National Chengchi University
Taipei, Taiwan 11623
chchen@nccu.edu.tw

Chung-Ching Tai
AI-ECON Research Center, Department of Economics, National Chengchi University
Taipei, Taiwan 11623
elliot@aiecon.org

Bin-Tzong Chie
AI-ECON Research Center, Department of Economics, National Chengchi University
Taipei, Taiwan 11623
chie@aiecon.org

Abstract In this chapter we conduct two experiments within an agent-based double auction market. These two experiments allow us to see the effect of learning and smartness on price dynamics and allocative efficiency. Our results are largely consistent with the stylized facts observed in experimental economics with human subjects. From the amelioration of price deviation and allocative efficiency, the effect of learning is vividly seen. However, smartness does not enhance market performance. In fact, the experiment with smarter agents (agents without a quote limit) results in a less stable price dynamics and lower allocative efficiency.

S.-H. Chen (ed.), *Genetic Algorithms and Genetic Programming in Computational Finance*

Keywords: Agent-Based Double Auction Markets, Genetic Programming, Quote Limits, Alpha Value, Allocative Efficiency

Introduction

In their seminal paper "**Allocative Efficiency of Market with Zero-Intelligence Trader**," Gode and Sunder posed an interesting question. *How much intelligence is required of an agent to achieve human-level trading performance?* The answer, as it first appeared, is a little surprising: *little*. How little? To make that message clearer, they called their agents *zero-intelligence* agents. These agents, when assigned a bargaining position in a standard double auction market, were simply bidding or asking randomly as if they had no capability to extract any useful information from the market. While these agents *individually* cannot bargain in an intelligent manner, their interactions via the market did *collectively* result in a near-100% allocative efficiency. They attributed this magic to Adam Smith's *invisible hand.*

> "Adam Smith's invisible hand may be more powerful than some have thought; it can generate *aggregate rationality* not only from individual rationality, but also from *individual irrationality.*" (Gode & Sunder, 1993, p.119, Italics added.).

It may be delicate to show that aggregate rationality does not rest upon individual rationality in some economic contexts. However, that does not include Gode and Sunder's case, and as a matter of fact, in their case it is *trivial* to show that allocative efficiency (aggregate rationality) does not rest upon individual rationality. To see this, let us simply assume that all agents are *truth-tellers.* In this regard, buyers would bid with their redemption values and sellers would ask with their unit costs. Obviously, allocative efficiency would then automatically be 100%.

Gode and Sunder did seem to assume that individual rationality would necessarily imply aggregate rationality (see the quotation above). They might also suppose that the allocative efficiency would be near 100% if all traders are *smart enough.* Consequently, their focus was to look for the *minimum* intelligence level which can generate aggregate rationality. What, however, may surprise them is that this minimum level is so low that you only need *dumb agents.*

Our picture is different: we do not assume that individual rationality necessarily implies aggregate rationality. Therefore, instead of the minimum level, we are looking at the other direction. We already have argued that the minimum level is not a problem at all: a group of innocent (honest) traders would result in a 100% allocative efficiency. What may make things uncertain is the case when traders are no longer innocent, but rather sophisticated. Hence, for us, the interesting question to ask is: *would smart or smarter traders reduce allocative efficiency?* As we shall see in this paper, the answer is positive.

In this paper we conduct agent-based simulations of DA markets with different intelligence levels of traders. Traders in one case are endowed with more space to act than traders in the other case. This extra space makes more sophisticated trading strategies possible to emerge, and hence can make traders in one case *potentially smarter* than traders in the other case.[1] Traders who are given this favor are called the *smart trader*, and traders who are not are called the *mediocre trader*. We then examine the allocative efficiency achieved within these two different setups. It is found that markets composed of mediocre traders realized 96% of the potential social surplus, whereas markets composed of smart traders realized only 88% of it.

Our finding for smart traders can be compared to the one Gode and Sunder saw for their zero-intelligence traders. While Gode and Sunder's finding shows that the invisible hand is more powerful than we thought, our finding shows the opposite. Moreover, to attain a higher allocative efficiency, the privilege given to traders should be deprived, and this deprivation can be interpreted as a kind of an intervention used to protect the market. Therefore, one way to summarize our study is as the following: *Smart agents do not necessarily bring goodness to the market. One purpose of regulation is to annihilate the evil-side of smartness.*

The rest of the paper is organized as follows. Section 1 shall briefly review the agent-based double auction markets used in this paper. Section 3 proposes experimental designs and measures of market performance. Section 4 presents and analyzes the simulation results, as Section 5 provides the concluding remarks.

1. Agent-Based Double Auction Markets: AIE-DA

Agent-based double auction (**DA**) markets were first seen in Dawid (1999), who applied the *single-population genetic algorithm* (**SGA**) to evolve traders' bids and asks, but not bargaining strategies per se. Chen (2000) proposed an agent-based DA market which is suitable for studying the evolution of bargaining strategies and which can be implemented with the software **AIE-DA**, developed by the **AI-ECON Research Center**. Chen (2001) and Chen, Chie, and Tai (2001) prepared a documentation accompanying this software, which is written in the language of **Delphi**, and is largely motivated by *object-oriented program-*

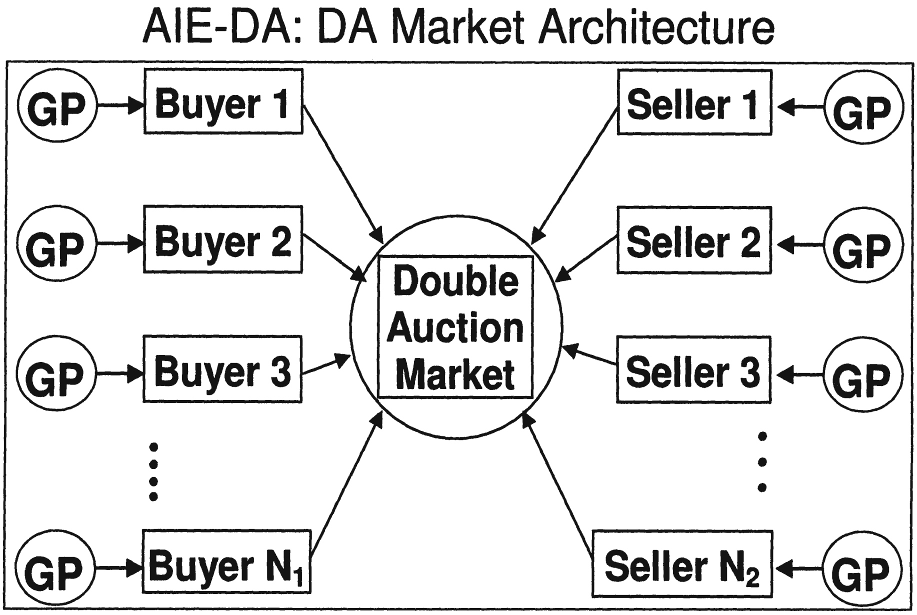

Figure 17.1. The AIE-DA Architecture: Multi-Population Genetic Programming.

ming (**OOP**). The experiments conducted in this paper will be based on this software.

All buyers and sellers in AIE-DA are *artificial adaptive agents* as described in Holland and Miller (1991). Each artificial adaptive agent is built upon *genetic programming.* The architecture of genetic programming used in AIE-DA is what is known as *multi-population genetic programming* (**MGP**). Briefly, we view or model an agent as *a population of bargaining strategies.*[2] Genetic programming is then applied to *evolving* each population of bargaining strategies. In this case, a society of bargaining agents consists of many populations of programs. This architecture is shown in Figure 17.1.

The evolution of bargaining strategies for each agent proceeds as follows. First, consider a counter called *generation.* At generation t, each agent is assigned K bargaining strategies, collectively denoted by Gen_t, whose determination will be explained later. For each trading period h, the agent takes a random strategy I as follows: $I \sim Uniform[1, K]$. At the end of the trading period, the profits of the chosen strategy i ($i \in [1, K]$) will be recorded as $\pi_{i,h}$. If strategy i is not chosen, then its profits will be counted as zero. After every H periods of trading, these k

strategies will be revised and renewed by standard genetic programming. The *fitness* is the mean profits:

$$\pi_i = \frac{\sum_{h=1}^{H} \pi_{i,h}}{\sum_{h=1}^{H} 1_{i,h}}, \tag{17.1}$$

where $1_{i,h}$ is 1 if i is selected at period h, otherwise it is 0.[3] The revision and renew process will generate a new generation of K strategies, and at this point, the counter shall move to generation $t+1$. The new generation of K strategies will be denoted by Gen_{t+1}. This revision and renew procedure is summarized as follows.

$$Gen_{t+1} \Longleftarrow \underbrace{[Reproduction] \bigvee [Crossover] \bigvee [Mutation]}_{Genetic\ \ Operators} \bullet (Gen_t) \tag{17.2}$$

As for the initial generation Gen_0, it will be generated by the *ramped half-and-half* method. This process will continue when t hits a prespecified number T. For the simulations in this paper, K is set to 50, H is 100, and T is 100.[4]

2. Experimental Designs

2.1. Quote Limit

To test the effect of *smartness* on allocative efficiency, a quote limit (upper bound for bid and lower bound for ask) is imposed in one experiment, but not the other. The purpose of a quote limit is to prevent buyers and sellers from bidding and asking an unprofitable price, and so a buyer cannot bid a price higher than his redemption value, and a seller cannot not ask a price lower than his unit cost. One may wonder why traders would be so *foolish* to bid or ask a price outside the quotation limit. Would they definitely make a loss? The answer is negative. Quoting a price outside the limit makes a trader's offer more lucrative, and enhances the chance of making a deal. Once the deal is won, depending on the trading mechanism, the trader may actually fulfill the transaction with a price which is different from his original quote.

To show an example, consider the **AURORA** computerized trading system developed by the Chicago Board of Trade. AURORA rules stipulate that only the holder of current bid (**CB**) or current ask (**CA**) is allowed to trade if CB $\geq$ CA.[5] By the AURORA rule, the actual transaction price (P) to fulfill a transaction will be somewhere between the current ask and the current bid. Let us assume the middle of them,[6]

Table 17.1. Values of Control Parameters for Genetic Programming

Number of Generations (T)	100
Population Size (K)	50
Evaluation Cycle (H)	100
Fitness Function	Mean Profits
Elitist Strategy	On
Number of Elites	1
Number of Strategies Chosen in Each Generation	100
Selection Scheme	Tournament
Tournament Size	5
Mutation Rate	0.05
Tree Mutation	0.1
Point Mutation	0.9
MaxDepth	17

i.e.,

$$P = \frac{CA + CB}{2}. \tag{17.3}$$

It is therefore clear that $CA \leq P \leq CB$, and so neither side would have to fulfill the transaction with their original quotes. This explains why the trader may take a risk of quoting a price outside the limit. This aggressive quotation should therefore be considered as a strategic behavior instead of a foolish one. Furthermore, this kind of aggressive bargaining strategy can actually emerges from the evolution of DA markets (Chen and Chie 2001).

2.2. Designs of Markets and Traders

Once the meaning of the quote limit is clear, we run two series of experiments. The first series of experiments are conducted without the quote limit, whereas the other series are conducted with it. This is essentially to say that we do not allow traders in the second series of experiments to evolve and develop aggressive bargaining strategies, while they are free to do so in the first series. By this limit, the "smartness" of the traders in the second series is restricted, whereas traders in the first are not. We therefore call traders in the first series *smart traders*, and traders in the second series *mediocre traders*.[7]

Twenty experiments were conducted for each series. In each experiment a token-value table, which is randomly generated, is applied to both

Table 17.2. List of Primitives, the Terminal Set and Function Set

Terminal Set	Function Set
Highest Token (**HT**) Next Token (**NT**) Lowest Token (**LT**)	Arithmetic Operators (+, −, ×, ÷)
Current Ask (**CASK**) Current Bid (**CBID**)	Absolute Value (**abs**)
Time Left Before the Termination of a Trading Period (**T1**) and Time Elaspe Since the Last Successful Trading (**T2**)	Logarithmic and Exponential Functions (**exp, log**)
The Average, Minimum and the Maximum Price of the Previous Trading Period (**PAvg, PMin, PMax**)	Trigonometric Functions (**sin, cos**)
The Average, Minimum and the Maximum Bid of Previous Trading Period (**PAvgBid, PMinBid, PMaxBid**)	Logical Operators (**If-Then-Else, If-Bigger-Then-Else**)
The Average, Minimum and the Maximum Ask of the Previous Trading Period (**PAvgAsk, PMinAsk, PMaxAsk**)	Extreme Operators (**max, min**)
Quiet (**Pass**)	Comparison Operator (>)
Ephemeral Random Constants (**C**)	

series. This table allows for four buyers and four sellers each with four units of tokens to trade. The 20 token-value tables (markets) generated are depicted in Figure 17.2. In each market, buyers and sellers' trading strategies evolved with genetic programming, whose control parameters are specified in Table 17.1 and 17.2.

2.3. Measures of Market Performance

In experimental economics, two measures are frequently used for market performance. One is the so-called *alpha value*, and the other is the *efficiency ratio.* The alpha value takes the *competitive equilibrium price* (p^*) as a benchmark, and attempts to measure the deviation degree between *market prices* and the *competitive equilibrium price.* More pre-

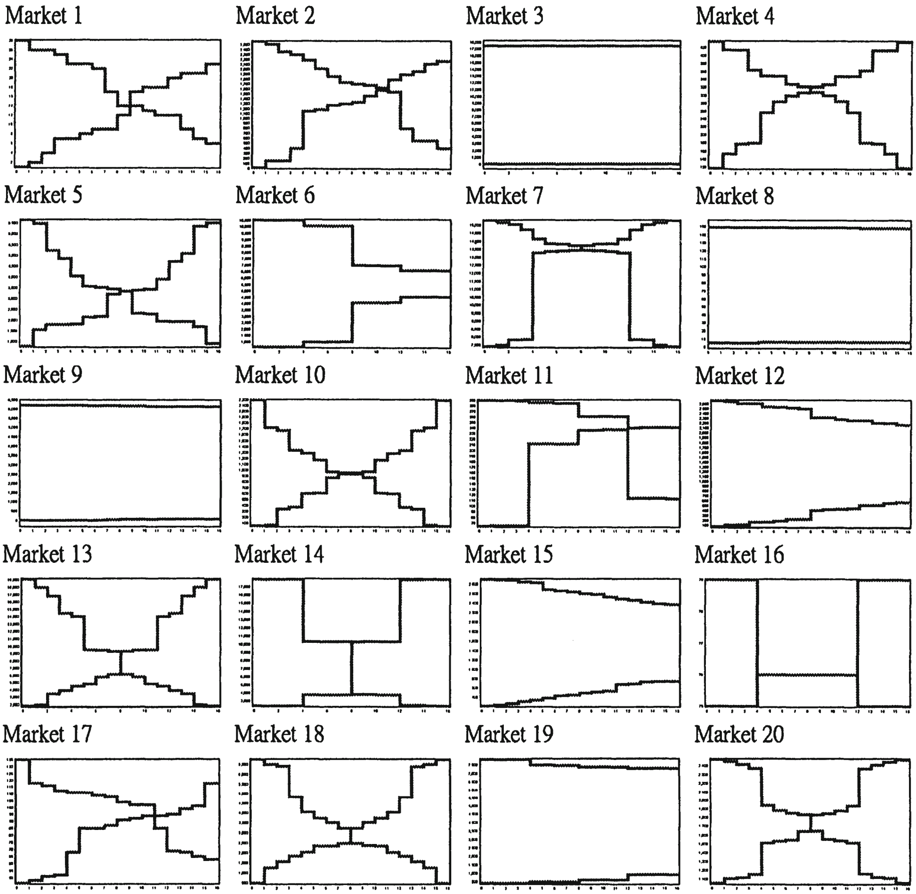

Figure 17.2. 20 Different Markets Used in the Experiments.

cisely, it is defined as follows,

$$\alpha = \frac{\sqrt{\frac{\sum_{i=1}^{n}(P_i - P^*)^2}{n}}}{P^*}, \tag{17.4}$$

where n is the number of transaction made in each trading period, and P_i is the market price of the ith transaction. One technical issue involved in this definition is the location of P^*. The meaning of competitive equilibrium price is clear when the demand and supply curves intersect at a single point, e.g., Markets 10 and 16 (See Figure 17.2). However, it is less clear when they intersect at an interval (Markets 7 and 20) or have no intersection at all (Markets 3 and 9). For the former case, we

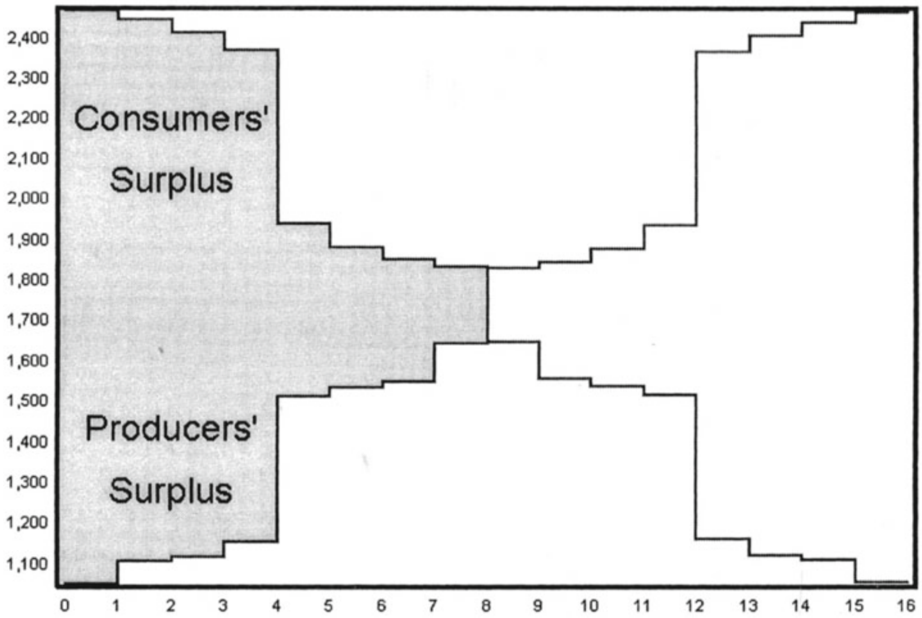

Figure 17.3. Consumers' Surplus & Producers' Surplus.

would define P^* as the midpoint of the intersecting interval, while for the latter case it is the midpoint between the lowest redemption value and the highest unit cost.

As to the second measure, we first take the sum of consumers' surplus (CS) and producers' surplus (PS) as the *potential social surplus* (See Figure 17.3). We then divide the realized surplus by the potential surplus, and post-multiply the quotient by 100%. The result is then a measure for allocative efficiency. It is mathematically described as follows. Let π_j be trader j's profits at a specific trading period,

$$\pi_j = \sum_{q \in bought} (V_{j,q} - P_q), \quad j \in Buyer, \tag{17.5}$$

and

$$\pi_j = \sum_{q \in sold} (P_q - V_{j,q}), \quad j \in Seller, \tag{17.6}$$

where $V_{j,q}$ is the redemption value of the qth unit for the jth buyer, or the unit cost of the qth unit for the jth seller. Term P_q is the transaction price of that unit. The *realized surplus* (RS) is then simply the sum of profits earned by all traders, namely,

$$RS = \sum_j \pi_j, \quad j \in traders. \tag{17.7}$$

The efficiency ratio β is

$$\beta = \frac{RS}{PS}, \tag{17.8}$$

where PS is the potential surplus. By this definition, $0 \leq \beta \leq 1$.

3. Experimental Results

3.1. Price Dynamics

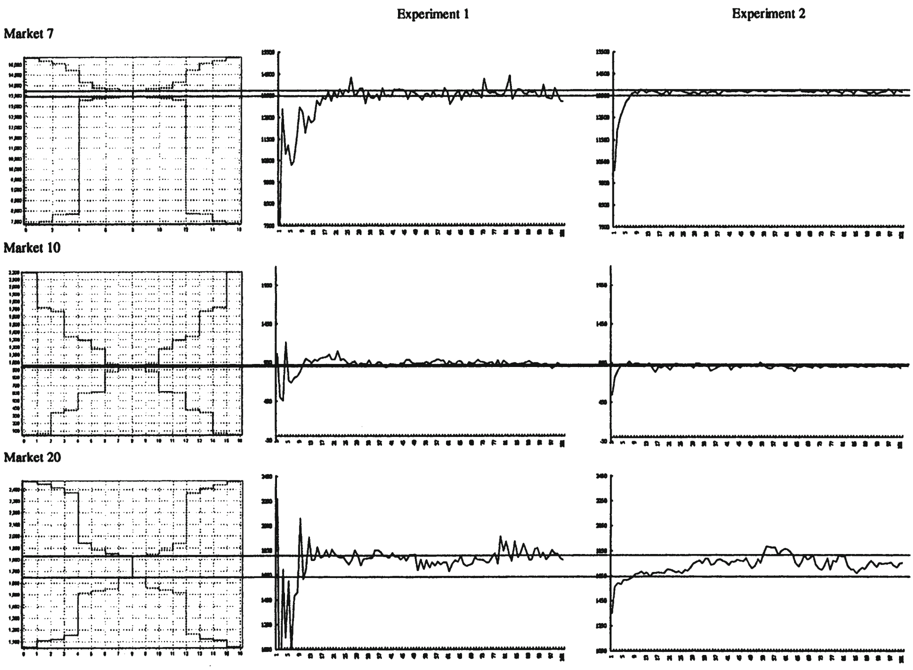

Figure 17.4. CASE 1: Time Series Plots of Price for Market 7, 10, and 20.

Our analysis of the simulation results will be based on the two measures introduced above. However, before proceeding to those measures, it will be useful to first look at the price dynamics of those markets. For this purpose, we demonstrate time series plots of the price for the following six markets: Markets 3, 7, 9, 10, 16, and 20 (Figures 17.4, 17.5, 17.6). For each figure, there are three plots. The leftmost plot is the market with its equilibrium price or equilibrium price interval. The middle plot is the time series of the price of experiment 1, whereas the rightmost plot is that of experiment 2. We shall distinguish these results by three separate cases. The first case refers to Markets 7, 10, and 20 (Figure 17.4). In these markets, we either have a unique equilibrium price (Market 10) or a tight equilibrium interval (Markets 7 and 20). Market prices in this case quickly move toward the equilibrium price (or price interval), and then slightly fluctuate around there. This result is basically consistent with what we learned from experimental economics with human subjects (Smith 1991).

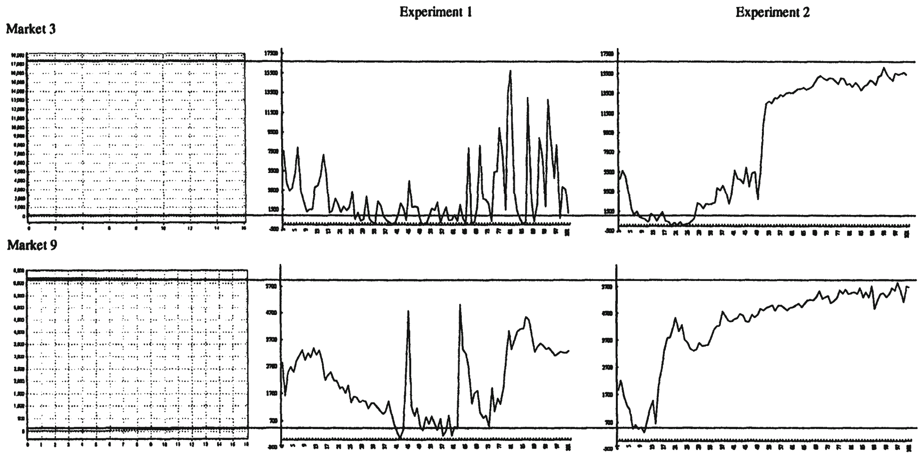

Figure 17.5. CASE 2: Time Series Plots of Price for Market 3 and 9.

The distinguishing feature of our DA markets actually starts from the second case, Markets 3 and 9 (Figure 17.5). The common characteristic of these two markets is that demand and supply curves are completely flat with no intersection. This case, while very intriguing, has almost been neglected in experimental economics literature.[8] How the price shall be determined under this circumstance is still an open question. Our simulation results from both markets seem to indicate that the price can shop around the whole interval, and is difficult to settle down to a narrow niche, but the price does not just randomly fluctuate. In fact, in Experiment 2, we see evidence of a slowly-upward moving trend. It begins with a price favorable to buyers, but then eventually turns to the seller side. For Experiment 1, Market 9 in particular, the price path is even more complicated. It starts with a downward trend, but then ends up with an upward trend. In the middle, the trend breaks occur several times.[9]

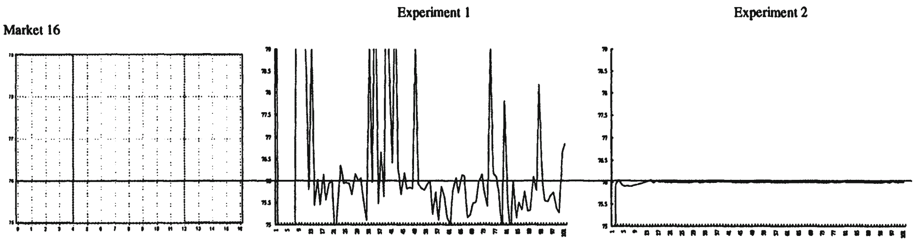

Figure 17.6. CASE 3: Time Series Plots of Price for Market 16.

We have so far not seen any qualitatively differences between Experiments 1 and 2. Their divergence appears to be clear in Case 3, Market 16 (Figure 17.6). In this market, demand and supply do intersect at a unique price. Based on our experience with Case 1, one might predict that the price will converge to this competitive equilibrium price; nevertheless, these two curves are completely flat before the intersection, and, as we have encountered in Case 2, the price may find it difficult to converge. The perplexity of this situation is also seen in our simulation results: Experiment 2 confirms the first conjecture, whereas Experiment 1 supports the second. Here is a good time to ask: *would a market with smart traders (a market without a quote limit) destabilize rather than stabilize the market?* We shall now deal with this question with more delicate statistics.

3.2. Alpha Value

As defined in Equation 17.4, the Alpha value measures the price deviation of a market period. Let $\alpha_{t,h}$ be the α value observed at the hth market period in the tth generation. In an ideal case, where $\lim_{t \to \infty} P_{t,h} \to P^*$,

$$\lim_{t \to \infty} \alpha_{t,h} = 0, \quad h = 1, ..., H. \tag{17.9}$$

The distribution (histogram) of $\{\alpha_{t,h}\}_{h=1}^{H}$ might also be expected to degenerate to the point zero. To see how far we are away from the ideal case, we draw the histogram of $\{\alpha_{t,h}\}_{h=1}^{H}$ for every 10 generations, i.e., t=0, 10,...100. To economize the presentation, the histogram is drawn

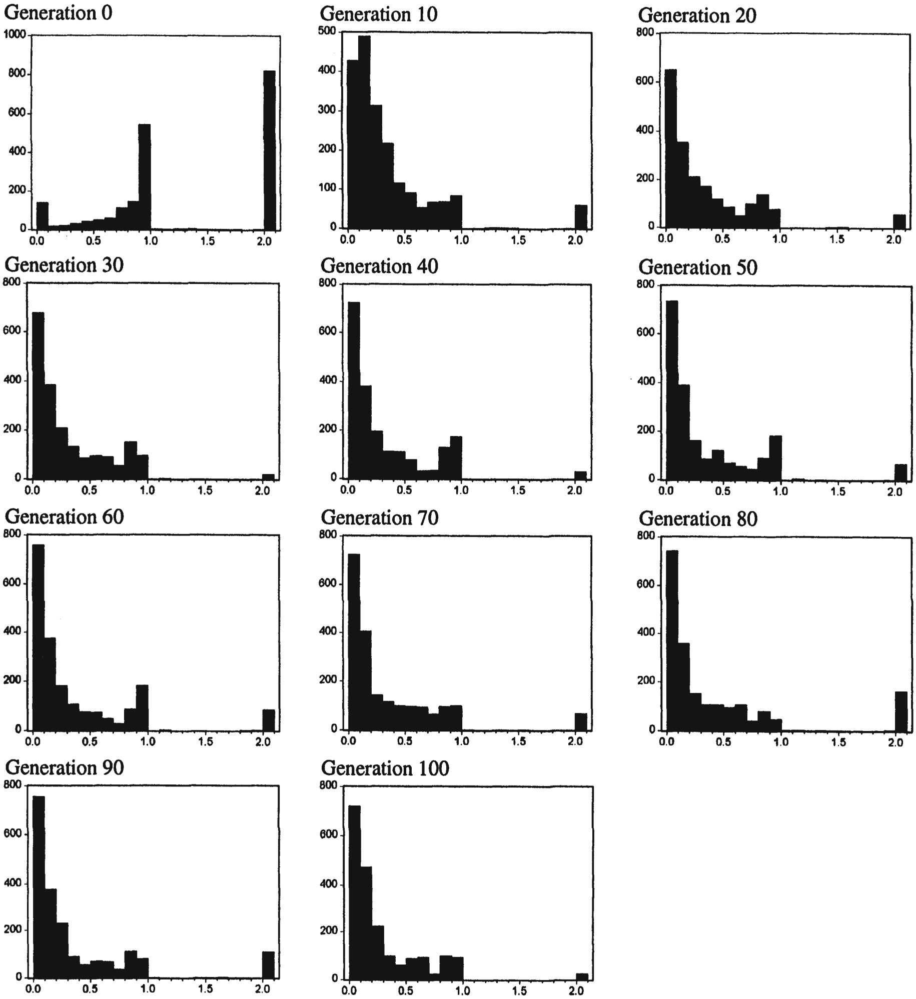

Figure 17.7. The Distribution of α Values: Experiment 1.

by pooling the $\{\alpha_{t,h}\}_{h=1}^{H}$ of all twenty markets. Figures 17.7 and 17.8 give the results of Experiments 1 and 2, respectively.[10]

These two figures roughly do indicate that these distributions exhibit a mass at zero and a skew to the right. Furthermore, Figure 17.9 plots the time series of the medium of $\{\alpha_{t,h}\}_{h=1}^{H}$ for both experiments. Here, we can see the effect of learning on price deviation: alpha values in both experiments tend to decrease in time. Drawing the two lines together also makes the effect of smartness transparent: the experiment with

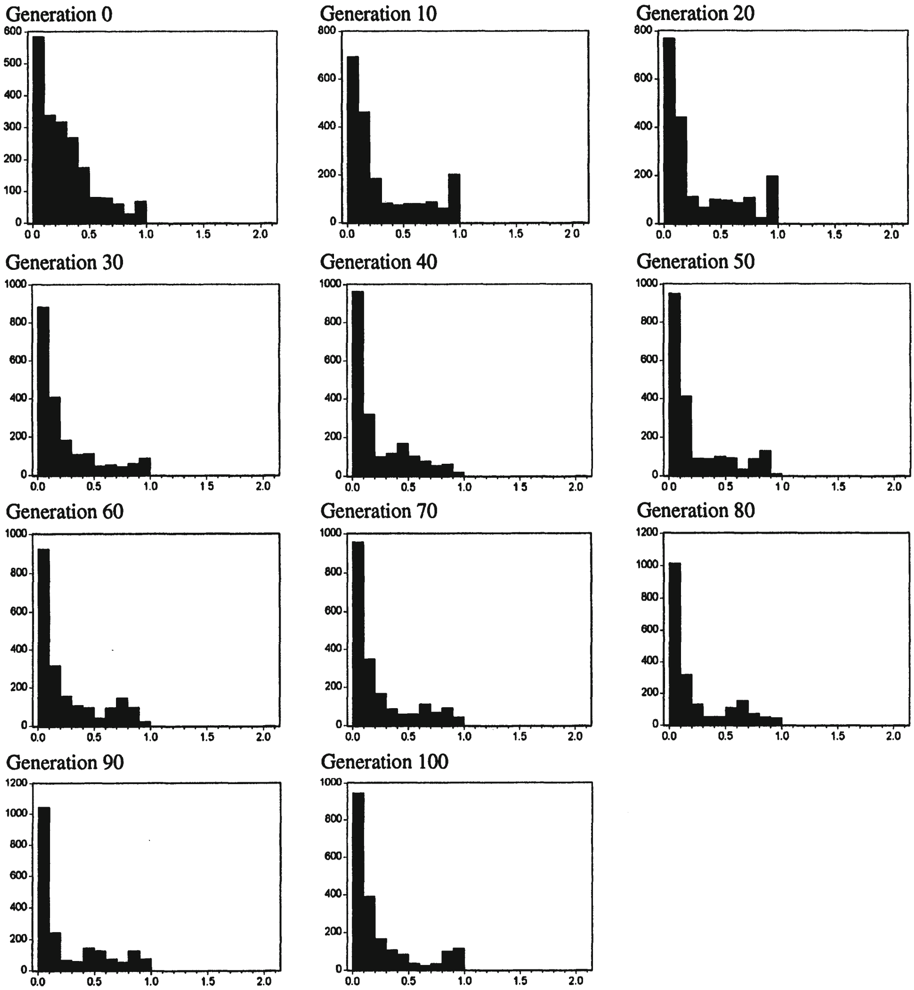

Figure 17.8. The Distribution of α Values: Experiment 2.

smarter agents (Experiment 1) tends to have a higher medium value at any point in time. This may help us answer the question posed above: *markets with smarter agents tend to be more fluctuating.* Hence, smarter agents play a *destabilizing* role for the market.

3.3. Beta Ratio

The β ratio, as defined in Equation 17.8, measures the allocative efficiency of a DA market. From the proceeding analysis above, we are

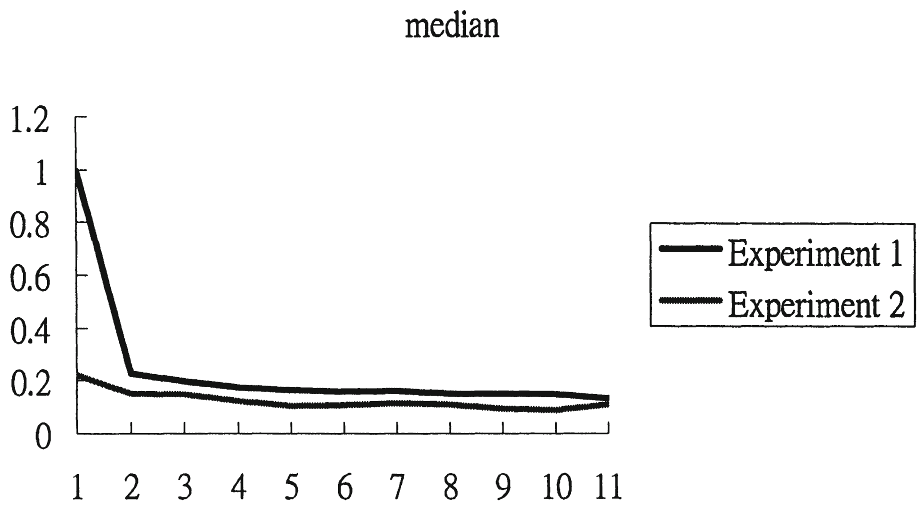

Figure 17.9. Time Series of the Median of the distribution of α Values.

interested in two things: first, the role of learning in allocative efficiency, and second, the role of smarter agents. Our first concern is shown in Figure 17.10. Figure 17.10 has three time series plots of the β ratio. The first time series plot shows that the β ratio increases with time. In this case, learning enhances the allocative efficiency, but, the second and the third do not. The third case even shows that allocative efficiency deteriorates after learning.

However, we have 20 markets for each experiment; therefore, we have numerous such time series plots. To economize our presentation, we use the *sup* function and the *inf* function to capture the essential characteristics of these three different time series plots.[11] Given an ordered series $\{x_t\}_{t=1}^{T}$, the *sup* function is defined as

$$sup_i(\{x_t\}_{t=1}^{T}) = \max\{x_t\}_{t=1}^{i}, \quad i = 1, ..., T, \tag{17.10}$$

and the *inf* function is[12]

$$inf_i(\{x_t\}_{t=1}^{T}) = \min\{x_t\}_{t=i}^{T}, \quad i = 1, ..., T. \tag{17.11}$$

The corresponding *sup* and *inf* functions of the three series are also drawn in Figure 17.10. As we can see from this diagram, if learning can enhance allocative efficiency, then these two functions should ideally display many upward jumps. On the contrary, if allocative efficiency deteriorates during the course of learning, then these two functions are

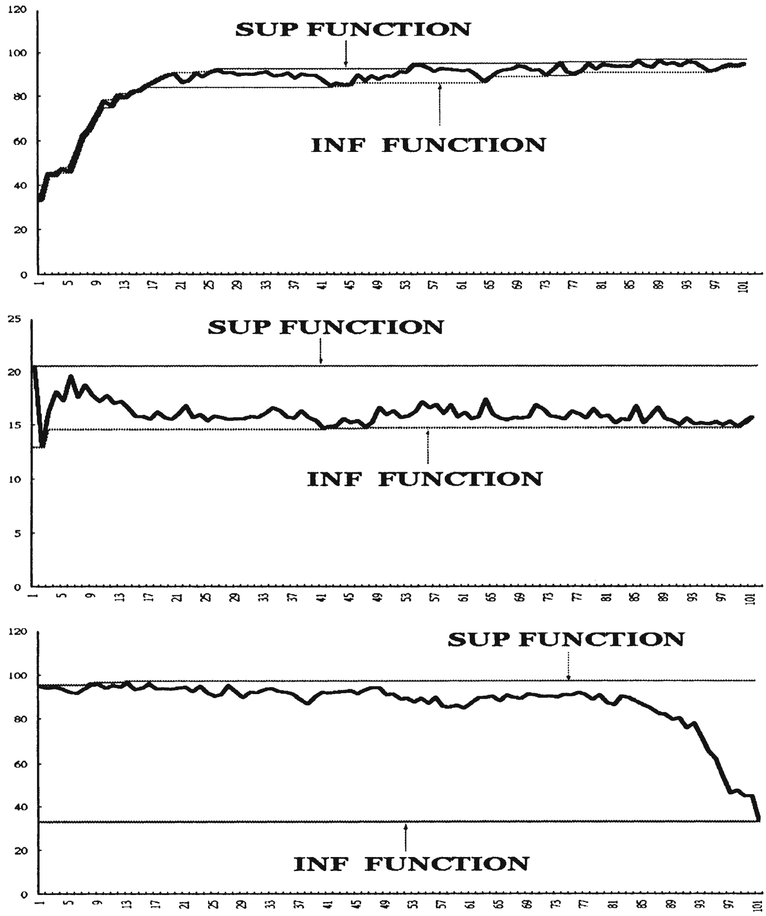

Figure 17.10. Three Possible Time Paths of Beta Ratio.

totally flat with no jumps at all. Therefore, by counting the number of jumps and the total jump size, one can effectively summarize the role of learning in allocative efficiency over different markets.

Figure 17.11 plots the number of jumps (on the y-axis) and jump size for the *inf* and *sup* function of the β values over 20 markets. Since each market may start with different initial ratios of β, we distinguish them by two different symbols, namely, diamond and box. Diamond stands for a lower initial ratio of β (below 0.4), whereas box stands for a higher initial ratio of β (between 0.4 and 0.6). Since jumps are prevalent for all markets, allocative efficiency becomes ameliorated as times goes on. Furthermore, for those markets with lower initial ratios of β, the

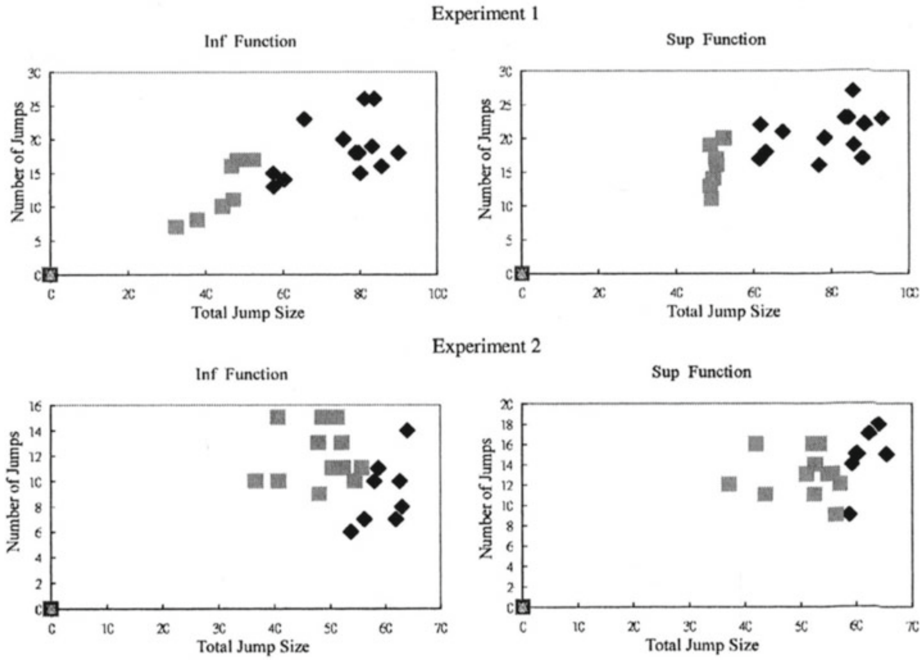

Figure 17.11. The Number of Jumps and Jump Size of Inf and Sup Functions.

amelioration degree is generally more significant than those with higher initial values.

To see the effect of smartness, the beta ratios of the initial generation and the last generation are shown in Table 17.3. The βs of Experiment 1 start with a range from 5% to 48%, and end up with a range from 79% to 98%. The βs of Experiment 2 are much higher: they start with a range from 33% to 59%, and end up with a range from 92% to 98%. If we compare these ratios pairwisely, then Experiment 1 is uniformly beaten by Experiment 2 in its resultant allocative efficiency, except for Market 9. Therefore, smarter traders do not enhance allocative efficiency.

4. Concluding Remarks

Our evidence is quite clear: smarter agents fail to enhance market performance.[13] They induce a relatively unstable price (higher alpha value) and lower allocative efficiency (lower β ratio). There is only thing left to address in this concluding section. *Why?*

Table 17.3. The Beta Ratio of the Initial and Last Generation

	Experiment 1		Experiment 2	
Market	Initial	End	Initial	End
1	29.09	94.67	56.93	97.56
2	32.09	85.60	44.01	96.14
3	45.56	83.44	41.19	97.00
4	8.59	90.91	45.32	93.37
5	11.11	92.28	33.75	96.78
6	46.66	89.28	46.74	94.73
7	6.41	86.45	59.69	96.29
8	48.32	96.63	46.72	97.25
9	46.09	98.31	38.50	97.23
10	9.92	79.33	42.12	94.41
11	33.42	90.91	55.26	96.00
12	45.50	92.62	39.66	97.66
13	4.47	80.07	34.59	98.59
14	8.57	84.27	33.78	96.34
15	48.42	95.18	36.86	98.78
16	6.44	89.19	46.69	98.00
17	34.14	89.10	43.62	98.11
18	3.80	89.37	38.82	92.55
19	47.26	79.40	38.02	94.26
20	5.46	89.03	45.80	94.51
Average		88.80		96.28

While it is not easy to provide a mathematical proof, we try to make our argument as plausible as possible. We shall build our argument based on the competition between *intra-marginal* and *extra-marginal* agents, who are the trading partners of *intra-marginal* and *extra-marginal* tokens. Intra-marginal and extra-marginal tokens are the tokens which are just inside or outside the equilibrium frontier. Without loss of generality, let us consider a very simple demand and supply schedule, as shown in Figure 17.12. Clearly, in this diagram only Seller 1 may have a successful trade. Let us also assume that Buyer 1 takes a random bid from the open interval (5,9). Given this bid and the imposition of the quote limit, Buyer 2 has no chance to beat Buyer 1 since his redemption value is only 5. In this case, Buyer 1 is an intra-marginal buyer, and Buyer 2 is the extra-marginal buyer.

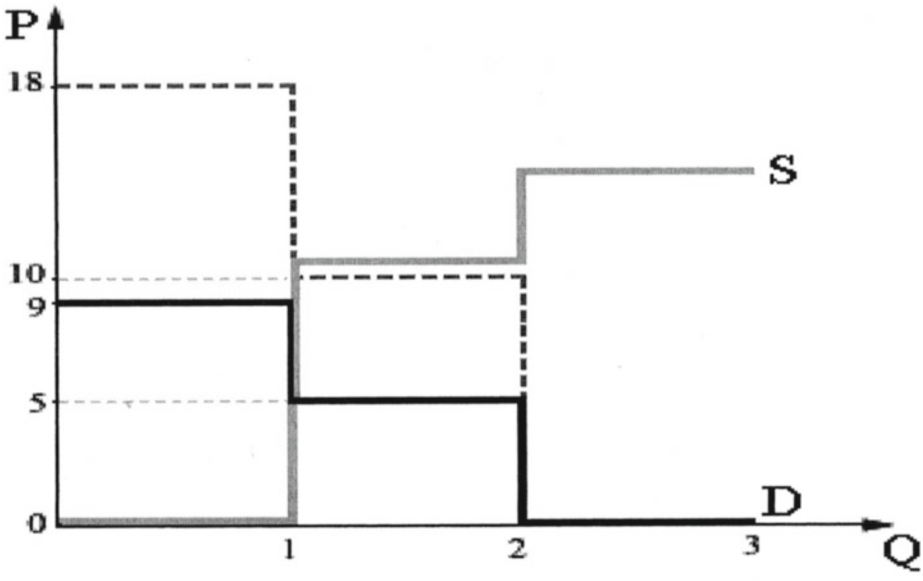

Figure 17.12. Removal of the Quote Limit and the Competition between Intra- and Extra-Marginal Agents.

Now consider the removal of the quote limit. Buyer 2 then *strategically* makes a bid up to 10 by not making a loss deal.[14] Supposing now that Buyer 1 still takes a random bid from (5, 9), the deal-winning chance for Buyer 2 will then jump from 0 to 0.2 if he takes a random bid from (5, 10). If Buyer 2 wins the deal, then the allocative efficiency derived becomes lower as opposed to the case where Buyer 1 wins the case. This gives us an explanation for why allocative efficiency will be deteriorated when the quote limit is removed.

One may next ask why Buyer 1 would not enlarge his bidding area, say up to 18. Yes, he will, if Buyer 2 continues to undertake the ambitious bidding. Nonetheless, when Buyer 2 is scared away, Buyer 1 may shift down his bidding area, and hence open the gate for the intruder again. Since we are using *multi-population genetic programming* to model our agents, it would be useful to make a distinction between *phenotype* and *genotype* in interpreting this dynamics. What we see from the outside is a competition between intra- and extra-marginal buyers, but what really happens inside (mentally) is the enduring competition between greedy strategies and cautious strategies. Greedy strategies nurse the growth of cautious strategies, which in turn do the same thing for the greedy strategies. The market just cannot settle its dynamics steadily. The in-and-out process is a generic property observed in many other agent-based markets.

Notes

1. Chen and Chie (2001) examined some *smart strategies* that evolved from their agent-based simulation of DA markets.
2. The number of bargaining strategies assigned to each bargaining agent is called the *population size*. AIE-DA Version 2 allows each agent to have at most 1000 bargaining strategies.
3. For the case when i is never selected, its mean profits would automatically be 0.
4. In other words, there totally are $H \times T = 10,000$ periods of trading for a single run of simulation.
5. Current bid refers to the highest bid at the current trading step, and current ask refers to the lowest ask. When CA is greater then CB, there shall be no match between buyers and sellers at the current step.
6. While the Aurora rules allow a random determination between CA and CB, we shall only consider the case by taking the average. See also Dawid (1999).
7. Notice that none of them in both series are born intelligent. They all have to learn and adapt to be intelligent. These terms only refer to the potentials which they may later develop.
8. To the best of our knowledge, Dawid (1999) is the only study that has drawn attention to this case.
9. Our results can be compared to Dawid (1999). Dawid (1999) found that price *always converges*, even though it did not necessary converge to the middle point of the demand and supply curves. However, there are several noticeable differences between our simulations and Dawid's. Firstly, Dawid did not use the Aurora Rule as the trading mechanism. Secondly, his market size is much bigger than us. Finally, his model of adaptive agents is also different from ours.
10. Since H is set to 100 in this paper, and there are 20 markets, there are 2,000 observations in each histogram.
11. This idea was first used in Chen, Kuo, and Lin (1996).
12. Actually, this is the inf function in reverse order.
13. Alternatively speaking, making everybody dumb by imposing a quote limit does not have a negative impact on market efficiency.
14. This upper limit is obtained by assuming that Seller 1 is a truth teller.

References

Chen, S.-H. (2000). "Toward an Agent-Based Computational Modeling of Bargaining Strategies in Double Auction Markets with Genetic Programming," in K.S. Leung, L.-W. Chan, and H Meng. (Eds.), *Intelligent Data Engineering and Automated Learning - IDEAL 2000: Data Mining, Financial Engineering, and Intelligent Agents*, Lecture Notes in Computer Sciences 1983, 517–531. Springer.

Chen, S.-H. (2001). "Evolving Bargaining Strategies with Genetic Programming: An Overview of AIE-DA Ver. 2, Part 1," in B. Verma, A. Namatame, X. Yao, H. Selvaraj, A. de Carvalho, and A. Ohuchi

(Eds.), *Proceedings of Fourth International Conference on Computational Intelligence and Multimedia Applications*, 48–54. IEEE Computer Society Press.

Chen, S.-H. and B.-T. Chie (2001). "The Schema Analysis of Emergent Bargaining Strategies in Agent-Based Double Auction Markets," in B. Verma, A. Namatame, X. Yao, H. Selvaraj, A. de Carvalho, and A. Ohuchi (Eds.), *Proceedings of Fourth International Conference on Computational Intelligence and Multimedia Applications*, 61–65. IEEE Computer Society Press.

Chen, S.-H., B.-T. Chie, and C.-C. Tai (2001). "Evolving Bargaining Strategies with Genetic Programming: An Overview of AIE-DA Ver. 2, Part 2," in B. Verma, A. Namatame, X. Yao, H. Selvaraj, A. de Carvalho, and A. Ohuchi (Eds.), *Proceedings of Fourth International Conference on Computational Intelligence and Multimedia Applications*, 55–60. IEEE Computer Society Press.

Chen, S.-H., J.-S. Kuo, and C.-C. Lin (1996). "From the Hayek Hypothesis to Animal Spirits: The Phase Transition based on Competitive Experimental Markets," in D. Vaz and K. Velupillai (Eds.), *Inflation, Institutions and Information: Essays in Honor of Axel Leijonhufvud*, 290–318. Macmillan.

Dawid, H. (1999). "On the Convergence of Genetic Learning in a Double Auction Market," *Journal of Economic Dynamics and Control*, 23, 1544–1567.

Gode, D. K. and S. Sunders (1993). "Allocative Efficiency of Market with Zero-Intelligence Trader: Market as a Partial Substitute for Individual Rationality," *Journal of Political Economy*, 101(1), 119–137.

Holland, J. H. and J. H. Miller (1991). "Artificial Adaptive Agents in Economic Theory," *American Economic Review: Papers and Proceedings*, 81(2), 365–370.

Smith, V. L. (1991). "Bidding and Auctioning Institutions: Experimental Results," *Papers in Experimental Economics*, 106–127, Cambridge: Cambridge University Press.