Market Mechanism Designs with Heterogeneous Trading Agents

Zengchang Qin
Berkeley Initiative in Soft Computing (BISC)
EECS Department, University of California
Berkeley, CA 94720, USA
zqin@cs.berkeley.edu

Abstract

Market mechanism design research is playing an important role in Computational Economics for resolving multi-agent allocation problems. A genetic algorithm was used to design auction mechanisms in order to automatically generate a desired market mechanism in agent based E-markets. In previous research, a hybrid market was studied, in which the probability that buyers rather than sellers are able to quote on a given time step, this probability was adapted by the GA which attempted to minimise Smith's coefficient of convergence. However, in previous experiments, all trading agents involved are of the same type or have identical preferences. This assumption does not hold in real-world markets which are always populated with heterogeneous agents. In this paper, the research of using evolutionary computing methods for auction designs is extended by using heterogeneous trading agents.

Keywords:
Genetic algorithm, multi-agent systems, intelligent agent, financial engineering, E-commerce, microstructure design

1 Introduction

In the first generation of E-commerce, bidders are generally humans who typically browse through well defined commodities with fixed prices via the Internet (e.g., Amazon.com). Just like the traditional marketplace, purchases are done with the prices made by sellers; buyers and sellers still have little freedom in transactions. For Customer-to-Customer (C2C) E-commerce like ebay.com, sellers and buyers are actually doing traditional trades, but through a new and more efficient medium. The freedom is also limited because both sellers and buyers still use the traditional methods of browsing to look for the goods they want. With the advent of agent technology, software agents can act as real-world traders in a virtual E-market. In comparison to human traders, such software agents have the advantages of being very fast, cheap and offer a tightly controlled environment in which a diverse range of experiments can be performed. A trading agent may represent a company or a customer hunting for maximized *utility* which means profit for the sellers or savings for the buyers. In such a way, freedom can be increased by allowing negotiation between opposite traders, i.e., sellers and buyers, in a large predefined cyberspace for transactions. As a result, commerce will become much more dynamic and the market less frictional. This kind of commerce is referred to as agent-mediated E-commerce or the second generation of E-commerce [5]. Since the traders will search in a very large space for matching their preferences, how to efficiently search the space and what protocols the traders have to follow in order to have a trustworthy and efficient market are all key problems for this new research area. In this paper, a method of using Genetic Algorithms (GAs) for trading protocol designs, or market mechanism designs, is discussed.

The market mechanism design is an important topic in Computational Economics and Finance [8]. By experimenting with zero-intelligence (ZI) agents, which simply generated random prices for bids or offers, Gode and Sunder [4] presented results that appear to indicate that a random guessing strategy can exhibit human-like behavior in Continuous Double Auction (CDA) markets. However, Cliff [1] indicated that the price convergence of ZI traders is predictable from a priori analysis of the statistics of the system, so that a more complex bargaining mechanisms or some "intelligence" is necessary for ZI traders. Consequently a type of agent with simple machine learning techniques was developed and referred to as zero intelligence plus (ZIP) agents. Further experiments by Das [3] showed that ZIP agents outperform their human counterparts.

By a series of experiments of exploring continuous auction space by ZIP agents via genetic algorithm, a hybrid auction with more desirable market dynamics, which had never been found in the real-world, was discovered

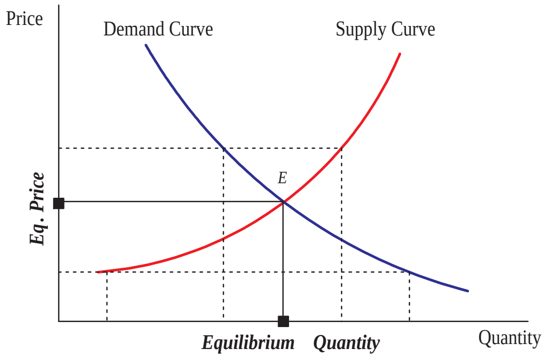

Figure 1. An schematic illustration of a supply-demand schedule, where the intersection E is the equilibrium.

[2]. However, in those experiments, the continuous auction space is not an exact analogue to real-world auctions. For example, single sided auctions such as English Auction (EA) and Dutch Flower Auction (DFA) in the proposed auction space are different from real-world EA and DFA. The question arises whether the hybrid auction is an artifact of the inexact auction model used in the experiments. Qin and Kovacs [10] proposed a new auction space that corrects the imperfectness of Cliff's model and yet still evolved the hybrid auction. In this paper, this research is continued by testing heterogeneous trading agents (a mixture of ZI and ZIP agents) in the marketplace evolved by a simple genetic algorithm.

2 Background Economics

In every classical economic model, demand and supply always play prominent roles. Supply is used to describe the quantity of a good or service that a household or firm would like to sell at a particular price. Demand is used to describe the quantity of a good or service that a household or firm chooses to buy at a given price. The intersection of the supply curve and demand curves is called the equilibrium, and the corresponding price and quantity are called, respectively, the *equilibrium price* and the *equilibrium quantity* (fig. 1). In case of prices beyond the equilibrium, the market will self-correct them to the equilibrium by an "invisible hand". At an equilibrium price, consumers get precisely the quantity of the good they are willing to buy at that price, and sellers sell out the quantity they are willing to sell at that price. Neither of them has any incentive to change. In a competitive market, the price actually paid and received in the market will tend to the equilibrium price. This is called the law of supply and demand [12].

In economics and game theory, interactions of traders consist of two components: a protocol and a strategy. Protocol defines the valid behavior of traders during the interaction. It is set by the marketplace owner and should be known publicly for all the participants. Strategy is privately designed by each agent to achieve their negotiation objectives within a protocol [6]. Moreover, the effectiveness of the strategy is very much dependent on the protocol: an optimal strategy for one protocol may perform very badly for other protocols. In a marketplace, the protocol is an"auction". It is the market mechanism by which buyers and sellers interact in this marketplace. Strategy is the adaptive behavior or "intelligence" of traders such as the ZIP agents' updating rules that will be discussed later.

There are many types of auctions. The following are some auctions used in this paper: English Auction (EA), sellers keep silent and buyers quote increasing bid-prices, and the buyer with highest bid is allowed to buy; Dutch Flower Auction (DFA), buyers keep silent and sellers quote decreasing offer-prices and the seller with lowest offer is allowed to sell. EA and DFA are also called single sided auctions because either buyers or sellers are active but not both. The Continuous Double Auction (CDA), one the most popular of all auctions, allows buyers and sellers to continuously update their bids/offers at any time in the trading period. The bids/offers are quoted simultaneously and asynchronously by buyers/sellers. At any time the sellers/buyers are free to accept the quoted bids/offers.

Classical economic theories always assume the number of traders in the market is infinite or very large. However, from a series of experiments performed over a six-year period starting in 1955, Smith [11] demonstrated that markets consisting of small numbers of traders could still exhibit equilibration to values predictable from classical microeconomic theory. This work helped to make a foundation for Experimental Economics [1].

In a given supply-demand schedule with n transactions between 'sellers' and 'buyers', the coefficient of convergence α ($0 \leq \alpha \leq 1$)is introduced to measure the deviation of transaction prices from the theoretical market equilibrium price p_0 [11]. α is calculated at the end based on transaction prices p_i for $i = 1, \cdots, n$. The coefficient of convergence is defined as follows:

$$\alpha = 100 \cdot \delta_0 / p_0 \tag{1}$$

where

$$\delta_0 = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(p_i - p_0)^2} \tag{2}$$

[1]Smith won the 2002 Nobel prize in Economics for his contributions in Experimental Economics. More information can be found at: http://nobelprize.org/economics/laureates/2002/

The E-market discussed in this paper as well as in [1, 2] and [10] is based on Smith's experiment and the α measure is used to evaluate the convergence of the market.

3 Automatic Auction Designs by Evolutionary Computing

Zero-intelligence (ZI) agents were initially proposed by Gode and Sunder [4] to explore the relationship between limited rationality, market institutions and the general equilibration of markets to the competitive equilibrium. Their fundamental discovery is that within the classical double auction market institution only the weakest elements of rationality (prices within a budget constraint are quoted randomly) need to be present for markets to exhibit high allocative efficiency and price convergence.

Zero intelligence plus (ZIP) agents, an augmented version of ZI agents with a simple machine learning technique, are fully described in [1]. Only a high-level description of the parameters for ZIP traders is given here: Each ZIP trader i is given a private secret limit price, λ_i, which for a seller is the price below which it must not sell and for a buyer is the price above which it must not buy (based on Smith's experiment). Each ZIP trader i maintains a time-varying profit margin $\mu_i(t)$ and generates quote-prices $p_i(t)$ at time t according to $p_i(t) = \lambda_i(1 + \mu_i(t))$ for sellers and $p_i(t) = \lambda_i(1 - \mu_i(t))$ for buyers. Trader i is given an initial value $\mu_i(0)$ (when $t = 0$) which is subsequently adapted over time using a simple machine learning technique known as the Widrow-Hoff (W-H) rule [7] which is well used in back-propagation neural networks. The W-H rule has a "learning rate" β_i that governs the speed of convergence between trader i's quote price $p_i(t)$ and the trader's idealised target price $\tau_i(t)$ which is determined by a stochastic function of last quote price with two small random absolute perturbations: $A_i(t)$ and $R_i(t)$. $A_i(t)$ is generated uniformly from the interval $[0, C_a]$ denoted by $\mathcal{U}[0, C_a]$ for sellers and $\mathcal{U}[-C_a, 0]$ for buyers. $R_i(t)$ is generated from $\mathcal{U}[1, 1 + C_r]$ for sellers and $\mathcal{U}[1 - C_r, 1]$ for buyers. Here C_a and C_r are called system constants. To smooth over noise in the learning, there is an additional "momentum" γ_i for each trader (momentum is also used in back propagation neural networks [7]).

For each ZIP agent i, its adaptation is governed by three real-valued parameters: learning rate β_i, momentum γ_i and initial profit margin $\mu_i(0)$. Because of the randomness and the uncertainty involved in trading, a trader's values for these parameters are assigned at initialization, using uniform distributions: for all traders, β_i is assigned a value at random from $\mathcal{U}(\beta_{min}, \beta_{min} + \beta_\Delta)$; and γ_i is from $\mathcal{U}(\gamma_{min}, \gamma_{min} + \gamma_\Delta)$ and $\mu_i(0)$ is from $\mathcal{U}(\mu_{min}, \mu_{min} + \mu_\Delta)$. Hence, to initialise an entire ZIP trader market it is necessary to specify values for the six market-initialisation parameters β_{min},β_Δ, γ_{min}, γ_Δ, μ_{min}, μ_Δ plus the other two system constants C_a and C_r. Clearly, any particular choice of values for these eight parameters can be represented as a vector:

$$V = [\beta_{min}, \beta_\Delta, \gamma_{min}, \gamma_\Delta, \mu_{min}, \mu_\Delta, C_a, C_r] \in \mathbf{R}^8$$

which corresponds to a single point in the 8-dimensional space of possible parameter values. A Genetic Algorithm can be used to explore this space for parameter optimization.

The fitness for each individual was calculated by monitoring price convergence in a series of n CDA market experiments, measured by weighting Smith's α measurement of convergence on the given supply-demand schedules. Each experiment lasted k "days" and the score of experiment number e is:

$$S(V_i, e) = \frac{1}{k} \sum_{d=1}^{k} w_d \alpha(d) \tag{3}$$

where $\alpha(d)$ is the value of α and w_d is the weight on the day d. According to the experiments in [2], all experiments last for 6 days and we place a greater emphasis on the early days of trading. The weights are set as follows: $w_1 = 1.75$, $w_2 = 1.50$, $w_3 = 1.25$ and w_4, w_5 and w_6 are all equal to 1.00. The fitness of the genotype V_i is evaluated by the mean score of n experiments:

$$F(V_i) = \frac{1}{n} \sum_{e=1}^{n} S(V_i, e) \tag{4}$$

Where $n = 50$ based on the empirical research in [9] which reported that the average of 50 independent runs of the trading experiments are fairly stable. The lower fitness a market has, the sooner the market approaches to the equilibrium and the smaller price variance the market has. GAs were used for optimising the parameters for ZIP agents and showed that evolved parameter settings via GAs perform significantly better than "educated guessing" in CDA and the same conclusion is also obtained in [9].

Now consider the case when we implement CDA. At time t, either a seller or a buyer will be selected to quote, which means that sellers and buyers have a fifty-fifty chance to quote. We use Q_s to denote the probability of the event that a seller offers. Then in CDA, $Q_s = 0.5$. For English Auction $Q_s = 0$ and Dutch Flower Auction $Q_s = 1$; which means, sellers cannot quote and sellers are always able to quote, respectively. The inventive step introduced in [2] was to consider the Q_s with values of 0.0, 0.5 and 1.0 not as three distinct market mechanisms, but rather as the two endpoints and the midpoint on a continuum referred as a continuous auction space. For other values, e.g., $Q_s = 0.1$, it can be interpreted as follows: on the average, for every ten quotes, there will be only one from sellers while 9 are from buyers. This also means, for a particular significant time t, the probability of a seller being the quoting trader is 0.1. The fact is, this kind of auction is never found in human-designed markets. However, no one knows whether this kind of *hybrid mechanism* in which $Q_s \neq 0, 0.5$ or 1.0 is preferable to human-designed ones. This motivates us to use a GA to explore with additional dimension Q_s ranging from 0 to 1 giving us the following genotype based on the old one by adding a new dimension Q_s:

$$[\beta_{min}, \beta_{\Delta}, \gamma_{min}, \gamma_{\Delta}, \mu_{min}, \mu_{\Delta}, C_a, C_r, Q_s] \in \mathbf{R^9}$$

According to the experiments in [2], the hybrid mechanisms are found to be the optimal auctions in 2 (SD_3 and SD_4, see section 5) of the 4 given schedules.

Although the case of $Q_s = 0.5$ is an exact approximation to the CDA in the real-world, the fact that a trader will accept a quote whenever the quoting price satisfies his expected price. For the two single sided extreme cases of $Q_s = 0.0$ and $Q_s = 1.0$, this model is not an exact analogue of the EA and DFA. Qin and Kovacs [10] proposed a more realistic auction space. All the following experiments are conducted in this realistic auction space. More detailed are available in [10].

4 Trading with Heterogeneous Agents in Continuous Double Auction

Smith's experiment (1962) qualitatively indicated that the relationship of the supply-demand schedule has an impact way in which transaction prices approached the equilibrium, even with a small number of participants, such a market would converge to equilibrium after only a small number of trading periods if the supply and demand remained constant. This experiment has been conducted by using ZI [4] and ZIP agents [2], respectively. In this paper, we will use a mixture of the same number of ZI and ZIP agents, which are referred to as heterogeneous agents experiments.

For all agents, the distribution of limit price determines the supply and demand curves for the experiment and their intersection indicates the theoretical equilibrium price and quantity. In the simulation of real marketplaces, we assume that each significant event (quoting, making deal or not making deal etc.) always occurs at a unique time. In the CDA market, at time t, an *active* trader (seller or buyer) i is chosen randomly to quote a price $p_i(t)$ to become the "current quote $q(t)$", where the active traders are ones who still have utility (goods or money) for deals. Next, all traders on the contra side (i.e. all buyers j if i is a seller, or all sellers j if i is a buyer) compare $q(t)$ to their current quote price $p_j(t)$ and if the quotes cross (i.e. if $p_j(t) \leq q(t)$ for sellers

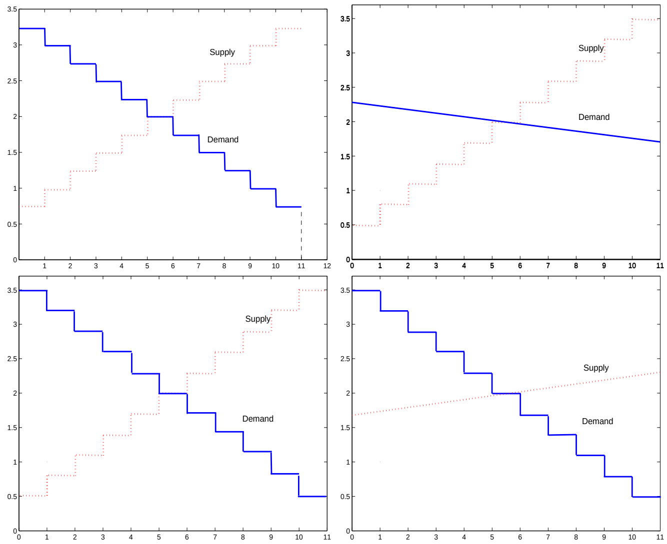

Figure 2. Supply-demand schedules for experiments: SD1, SD2 (upper) and SD3, SD4 (bottom).

or $p_j(t) \geq q(t)$ for buyers) then the trader j is able to accept. If no traders are able to accept, the quote is regarded as "ignored". For ZIP traders, either the current quote is accepted or ignored and the traders update their profit margins $\mu(t)$ using the W-H rule. For example, suppose the last quote is an offer and was accepted at price q then any sellers for which their price is less than q should raise their profit margin with learning rate of β_i. The details about the updating rules for ZIP agents can be found in [1]. For ZI traders, the previous transaction prices and the status of the last offer do not have cause any influence on their further actions (ZI traders are not intelligent, they only quote prices randomly).

5 Experimental Studies

In this section, we conduct a series of experiments of evolutionary designs of market mechanism based on heterogeneous agents where ZI and ZIP agents have the approximately same number. The auction space model is the one proposed in [10]. All experiments are based on four given supply-demand schedules: SD_1, SD_2, SD_3 and SD_4 (see figure 2). There are 22 trading agents in the experiments, 11 sellers and 11 buyers, each of them is initialized with one unit of goods and their limit prices are distributed as supply and demand curves show. The vertical axis represents price and the equilibrium price is 2.00 for all these 4 given schedules. Each schedule of supply and demand curves is stepped. This is because the commodity is dealt in indivisible discrete units, and there are only a small number of units available in the market. Thus, supply and demand in this simple market differs appreciably from the smoothly sloping curves of an idealised market. These are the same schedules have also been used in previous studies. It is convenient for us to do comparison studies.

In the market evolution experiments, a simple GA is used to minimize the fitness value (see equation 4) given 25 independent runs of trading experiments. The values for key parameters of GA are given in table 1. The Genetic Algorithm (GA) has become a standard optimization technology and there are many good literatures available for the readers who are not that familiar with it. Here, we only give the parameter initialization of the standard GA without a general introduction of it. Population size is 20 and each parameter is coded with 8 bits, crossover rate is a constant with the value of 0.7 and mutation rate is 0.015. Elitism strategy is applied which means that the fittest individual in each generation is logged. We run 600 generations in a single experiment. However, one of the drawbacks of using a GA is that it cannot be guaranteed that the solution on which the population eventually converges is a global rather than a local optimum. Thus we gain formal simplicity at the cost of computation. We run the entire process of evolution many times independently and reduce the effect of mutation as time goes by, to encourage convergence. The results of Q_s represented here are based on 25 independent runs of the GA on the given 4 supply-demand schedules and the average results with standard deviation through generation 600 are shown in figure 3.

As we can see from the figures, although Q_s values converges to real-world auctions in 3 of the 4 given schedules, we still found a hybrid auction in SD_4. Comparing the ZIP agents in the old auction space and the new auction space, the only difference is SD_3. Both in the old auction (Cliff) and new auction space (Qin and Kovacs) with ZIP agents, there were hybrid auctions found by GAs. Cliff [2] presented a result of using only ZI agents given SD_3 and the hybrid auction was found. However, the Q_s values for these hybrid auctions are different: $Q_s = 0.39$ for experiments with ZI agents only, $Q_s = 0.16$ for ZIP agents in the old auction space and $Q_s = 0.23$ for ZIP agents in the new auction space [10]. Here in the experiment with heterogeneous agents which are a mixture of ZI and ZIP agents, the optimal auction is CDA but not a hybrid one. We believe that the optimal auction for a market is related to the supply-demand schedule given. So far, we just demonstrated with empirical studies due to the complexity of such problems. The theoretic relations among hybrid auction, supply-demand schedule, the number of agents and other factors are considered as a future work. However, we demonstrated that given a particular supply-demand schedule, we can use some machine learning technology to find the optimal auction for such a market.

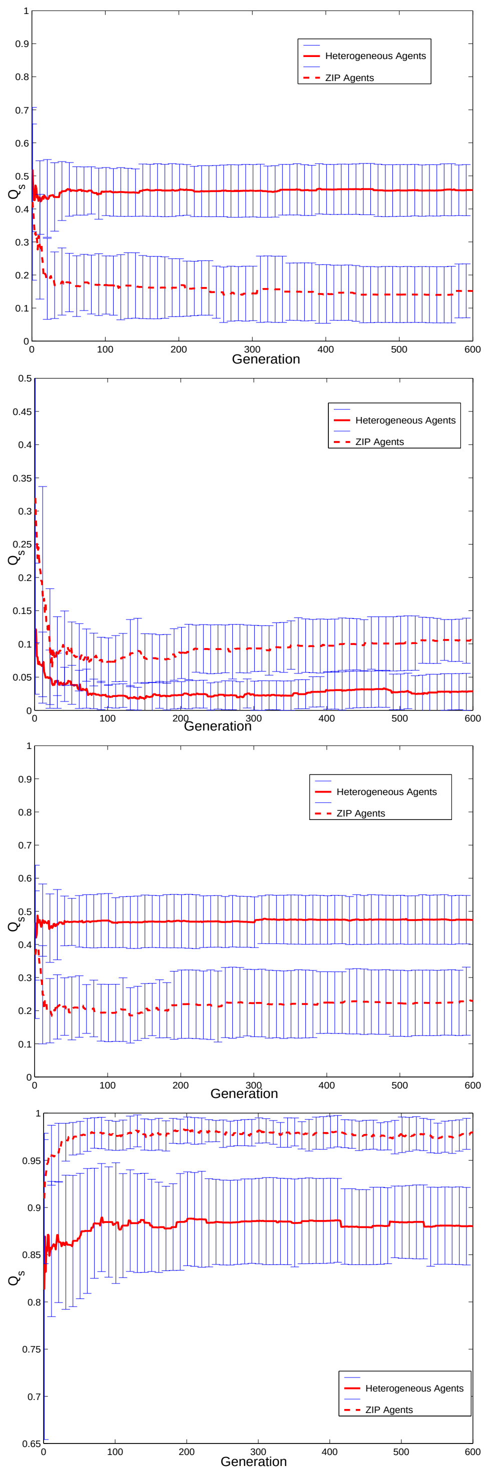

Figure 3. The comparisons of evolutionary trials of Q_s for ZIP (dot lines) and heterogeneous agents (solid lines) on schedules SD1 to SD4 through 600 generations.

Table 1. Parameter settings for the GAs used in the evolving market mechanism experiments.

Parameters for GA	*Value*	*Parameters for GA*	*Value*
Population	20	Num. of Parameters	9
Maximum Generation	600	Elitism	YES
Crossover Rate	0.7	Mutation Rate	0.015
Bits per Parameter	8	Selection Method	Rank Selection

6 Conclusions and Discussions

In this paper, we reviewed the method of using genetic algorithms for designing market mechanism and redid the experiments of [10] with heterogeneous agents instead of one sort of trading agent. Based on the evolving market mechanism experiments with heterogeneous agents, we found that in the 3 of the 4 given supply-demand schedules, the Q_s values converge to real-world auctions such as CDA and EA. In the last schedule SD_4, the auction with the most desired market dynamics was a hybrid mechanism.

We would like to point out that this is not a trivial academic point: although the efficiency of automatically designed markets is only a few percentage better than those designed by human, the economic consequences could be highly significant. For example, the total vaue of trades on the CDA-based New Yourk Stock Exchange (NYSE) for the year 2000 was 110060 billion US dollars and if only 0.1% of the liquidity could be saved by using a more efficient market, the value would be about 10 billion dollars. Although the experiments presented in this paper are far from real-world applications. From the academic point of view, we provide a new way of designing microstructure of markets.

Acknowledgment

The author thank Tim Kovacs, Dave Cliff and Andrew Byde for some useful comments in the early stage of this work. This work was conducted when the author was at the Bristol University and partly funded by ORS UK and BT/BISC Fellowship.

References

[1] D. Cliff, Minimal-intelligence agents for bargaining behaviors in market-based environments. *Technical Report HPL-97-91*, Hewlett-Packard Laboratories, 1997.

[2] D. Cliff, Explorations in evolutionary design of online auction market mechanism, *Electronic Commerce Research and Applications*, Vol. 2, pp. 162-175, 2003.

[3] R. Das, J. E. Hanson, J. O. Kephart, and G. Tesauro, Agent-human interactions in the continuous double auction, *Proceedings of the IJCAI-2001*, Seattle, USA, 2001.

[4] D. K. Gode and S. Sunder, Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality, *Journal of Political Economy*, 101(1), 119-137, 1993.

[5] M. He, N. Jennings, and H. F. Leung, On agent-mediated electronic commerce, *IEEE Trans. on Knowledge and Data Engineering*, Vol. 15, No. 4, 2003.

[6] M. He, H. F. Leung, and N. Jennings, A fuzzy logic based bidding strategy for autonomous agents in continuous double auctions, *IEEE Trans. Knowledge and Data Engineering*, 15(6), pp. 1345-1363, 2003.

[7] J. Hertz, A. Krogh and R. G. Palmer, *Introduction to The Theory of Neural Computation*, Addison-Wesley, 1991.

[8] B. LeBaron, Agent-based computational finance: Suggested readings and early research, *Journal of Economic Dynamics and Control*, Vol. 24, pp. 679-702, 2000.

[9] Z. Qin, *Evolving Marketplace Designs by Artificial Agents*, MSc Dissertation, Computer Science, University of Bristol, 2002.

[10] Z. Qin and T. Kovacs, Evolution of realistic auctions, M. Withall and C. Hinde Ed. *Proceedings of the 2004 UK workshop on Computational Intelligence*, pp. 43-50, Loughborough, UK, 2004.

[11] V. L. Smith, An experimental study of competitive market behavior, *Journal of Political Economy* 70, pp. 111-137, 1962.

[12] J. E. Stiglitz and J. Driffill, *Economics*, Chapter 4: Demand, supply and price, W. W. Norton & Company, Inc., 2000.