ELSEVIER

Journal of Economic Dynamics & Control
25 (2001) 363–393

JOURNAL OF Economic Dynamics & Control

www.elsevier.nl/locate/econbase

Evolving traders and the business school with genetic programming: A new architecture of the agent-based artificial stock market☆

Shu-Heng Chen[a,]*, Chia-Hsuan Yeh[b]

[a]*AI-ECON Research Group, Department of Economics, National Chengchi University, Taipei, 11623 Taiwan*

[b]*AI-ECON Research Group, Department of Finance I-Shou University, Kaohsiung County, 84008 Taiwan*

Accepted 22 February 2000

Abstract

In this paper, we propose a new architecture to study artificial stock markets. This architecture rests on a *mechanism* called 'school' which is a *procedure* to map the phenotype to the genotype or, in plain English, to uncover the secret of success. We propose an *agent-based model of 'school'*, and consider school as an evolving population driven by single-population GP (SGP). The architecture also takes into consideration traders' search behavior. By *simulated annealing*, traders' search density can be connected to *psychological factors*, such as *peer pressure* or economic factors such as the *standard of living*. This market architecture was then implemented in a standard artificial stock market. Our econometric study of the resultant artificial time series evidences that the return series is independently and identically distributed (iid), and hence supports the efficient market hypothesis (EMH). What is interesting though is that this iid series was generated by traders, who do not believe in the *EMH* at all. In fact, our study indicates

☆The paper is adapted from a version presented at the *Seventh Annual Symposium of the Society for Nonlinear Dynamics and Econometrics* held in New York on March 18–19, 1999, and the *Fifth International Conference of the Society for Computational Economics* held in Boston College on June 24–26, 1999. The authors are grateful for the comments from the participants of these two conferences and the constructive suggestion from one anonymous referee. Research support from NSC grants No.88-2415-H-004-008 and No. 87-I-2911-15 is also gratefully acknowledged.

* Corresponding author. Tel.: 886-2-29387308; fax: 886-2-27386874.

E-mail addresses: chchen@nccu.edu.tw (S.-H. Chen), spockyeh@ms38.hinet.ne (C.-H. Yeh).

PII: S 0 1 6 5 - 1 8 8 9 (0 0) 0 0 0 3 0 - 0

that many of our traders were able to find *useful signals* quite often from business school, even though these signals were short-lived.

JEL classification: G12; G14; D83

Keywords: Agent-based computational economics; Social learning; Genetic programming; Business school; Artificial stock markets

1. Background and motivation

Over the past few years, genetic algorithms (GAs) as well as genetic programming have gradually become a major tool in *agent-based computational economics* (ABCE). According to Holland and Miller (1991), there are two styles of GAs or GP in ABCE, namely, *single-population GAs/GP* (SGA/SGP) and *multi-population GAs* (*GP*) (MGA/MGP). SGA/SGP represents each agent as a *chromosome* or a *tree*, and the whole population of chromosomes and trees are treated as a society of market participants or game players. The evolution of this society can then be implemented by running canonical GAs/GP. Arifovic (1995, 1996), Miller (1996), Vila (1997), Arifovic et al. (1997), Bullard and Duffy (1998a,b, 1999), Staudinger (1998) are examples of SGA, while Andrews and Prager (1994), Chen and Yeh (1996, 1997, 1998), and Chen et al. (1996) are examples of SGP. MGA/MGP, in contrast, represents each agent as *a society of minds* (Minsky, 1986). Therefore, GAs or GP is actually run inside each agent. Since, in most applications, direct conversations (imitations) among agents do not exist, this version of applications should not be mistaken as the applications of *parallel and distributed GAs/GP*, where communications among '*islands*' do exist. Examples of MGA can be found in Palmer et al. (1994), Tayler (1995), Arthur et al. (1997), Price (1997), Heymann et al. (1998).

While these two styles of GAs/GP may not be much different in engineering applications, they do answer differently for the fundamental issue: '*who learns what from whom?*' (Herreiner, 1998). First, agents in the SGA/SGP architecture usually learn from other agents' experiences, whereas agents in the MGA/MGP architecture only learn from their own experience. Second, agents' interactions in the SGA/SGP architecture are direct and through *imitation*, while agents' interactions in the MGA/MGP architecture are indirect and are mainly through *meditation*. It is due to this difference that SGA/SGP is also called *social learning* and MGA/MGP *individual learning* (Vriend, 1998). At the current state, the SGA/SGP architecture is much more popular than the MGA/MGP architecture in ABCE.

In addition to its easy implementation, the reason for the dominance of SGA/SGP in ABCE is that economists would like to see their genetic operators (reproduction, crossover, and mutation) implemented within a framework of social learning so that the population dynamics delivered by these genetic operators can be directly interpreted as market dynamics. In particular, some interesting processes, such as imitation, 'following the herd', rumors dissemination, can be more effectively encapsulated into the SGA/SGP architecture.

However, it has been recently questioned by many economists whether SGA/SGP can represent a sensible learning process at all. One of the main criticisms is given by Harrald (1998), who pointed out the traditional distinction between the *phenotype* and *genotype* in biology and doubted whether the adaptation can be directly operated on the genotype via the phenotype in social processes. Back to Herreiner's issue, if we assume that agents only imitate others' *actions* (phenotype) without adopting their *strategies* (genotype), then SGA/SGP may be immune from Harrald's criticism. However, imitating other agents' actions are a very minor part of agents' interactions. In many situations, such as financial markets and prisoners' dilemma games, it would be hopeless to evolve any interesting agents if they are assumed to be able to learn only to '*buy and hold*' or '*cooperate and defect*'. More importantly, what concerns us is *how they learn the strategies behind these actions*. But, unless we also assume that strategies are *observable*, it would be difficult to expect that they are *imitable*. Unfortunately, in reality, strategies are in general not observable. For instance, it is very difficult to know the forecasting models used by traders in financial markets. To some extent, they are *secrets*. What is observable is, instead, only a sequence of trading actions. Therefore, Harrald's criticism is in effect challenging all serious applications of SGA/SGP in ABCE.

Although Harrald's criticism is well acknowledged, we have seen no solution proposed to tackle this issue yet. At this stage, the only alternative offered is MGA/MGP. In fact, it is interesting to note that many applications which heavily rely on evolution operated on the genotype (strategies) tend to use MGA/MGP. Modeling financial agents is a case in point. What is ironic is that this type of application is in essence dealing with human interaction and thus requires an explicit modeling of *imitation*, *speculation* and *herd behavior*. As a result, MGA/MGP is not really a satisfactory response to Harrald's criticism.

In this paper, we plan to propose a new architecture and hence a solution to Harrald's criticism. This architecture rests on a *missing mechanism*, which we think is a key to Harrald's criticism. The missing mechanism is what we call '*school*'. Why 'school'? To answer Harrald's criticism, one must resolve the issue '*how can unobservable strategies be actually imitable*'? The point is *how*. Therefore, by the question, what is missing in SGA/SGP is a *function* to show how, and that function is what we call 'school'. Here, 'school' is treated as a *procedure*, a procedure to map the phenotype to the genotype, or in plain English, to uncover the secret of success. This notion of 'school' goes well with what school usually means in our mind. However, it covers more. It can be mass media, national library, information suppliers, and so on. Warren Buffett may not be generous enough to share his secrets of acquiring wealth, but there are hundreds of books and consultants that would be more than happy to do this for us. All these kinds of activities are called 'schooling'. Therefore, if we supplement SGA/SGP with a function 'school', then Harrald's criticism can, in principle, be solved.

Nevertheless, to add 'school' to an evolving population is not that obvious. Based on our earlier description, 'school' is expected to be a collection of most updated studies about the evolving population (evolving market participants). So, to achieve this goal, 'school' itself has to evolve. The question is how? In this paper, we propose an *agent-based model of 'school'*. More precisely, we consider school as an evolving population driven by single-population GP (SGP). In other words, 'school' mainly consists of faculty members (agents) who are competing with each other to survive (get tenure or research grants), and hence the survival of the fittest principle is employed to drive the evolution of faculty the way it drives the evolution of market participants. To survive well, a faculty member must do her best to answer *what is the key to success in the evolving market*. Of course, as the market evolves, the answer also needs to be revised and updated.

Once 'school' is constructed with the agent-based market, the SGP used to evolve the market is now also run in the context of school. The advantage of this setup is that, while the SGP used to evolve the market suffers from Harrald's criticism, the SGP used to evolve 'school' does not. The reason is simple. To be a successful member, one must publish as much as she knows and cannot keep anything secret. In this case, observability and imitability (replicatability) is not an assumption but a rule. In other words, there is no distinction between the genotype and phenotype in 'school'. Hence, Harrald's criticism does not apply and SGP can be 'safely' used to evolve 'school'.

Now, what happens to the original SGP used to evolve the market? This brings up the second advantage of our approach. Since the function of school is to keep track of strategies (genotypes) of market participants and to continuously generate new and promising ones, any agent who has pressure to imitate other agents' strategies or to look for even better strategies can now just consult 'school' and see whether she has any good luck to have a rewarding search. So, the original operation of SGP in the market can now be replaced by SGP in 'school' and a *search* procedure driven by the survival pressure of agents. Agents can still have interaction on the phenotype in the market, but their interaction on the genotype is now indirectly operated in 'school'.

An interesting aspect of this approach is to explicitly model the interaction between '*school*' and *the market* by introducing a co-evolution model. To survive, school must adapt to market dynamics. On the other hand, market dynamics generate *students* for 'school' who, in turn, bring the knowledge learned from 'school' back to the market, and that knowledge may have further impact on market dynamics. While agent-based modeling is a *bottom-up* approach, one may use a system of two nonlinear difference equations, governing the dynamics of 'school' and the market, as a *top-down* 'summary'.

The difference between our proposed architecture and SGA/SGP and MGA/MGP is also illustrated in Figs. 1–3. Fig. 1 depicts the market architecture represented by SGA/SGP. The top of Fig. 1 is the *market* as a single object, and the bottom is a *population of directly interacting heterogeneous agents*. The *direct interaction* is characterized by the symbol '$\leftrightarrow$' among them. By this architecture, the information (knowledge) about the market is *openly distributed* among all agents. Nothing is kept secret. In between is a symbol '$=$' (equivalent to), which means that market dynamics is equivalent to the evolution of this population of directly interacting agents.

Fig. 2 gives the market architecture represented by multi-population GP. The market remains at the top, but there are two essential differences as opposed to the previous figure. First, the single symbol $=$ is replaced by a series of $\Leftrightarrow$s. Under these $\Leftrightarrow$s is a population of *indirectly* interacting agents. By 'indirectly', we mean that these agents are interacting only through a bulletin board. Imagine that each agent sits in her office and watches the world from the web. They have no direct contact with one other, physically, and in some sense, mentally as well. The information (knowledge) about the market is now *privately distributed* among agents. Each agent has her own world and keeps her own secrets.

The point here is that other agents' minds are not directly observable, and hence not imitable. Within each agent's mind, there is a society of minds. The evolution of this society is driven by GP. Within this architecture, agents basically learn from her own experience, and not from other agents' experiences. Thus, it is a typical model of *individual learning*.

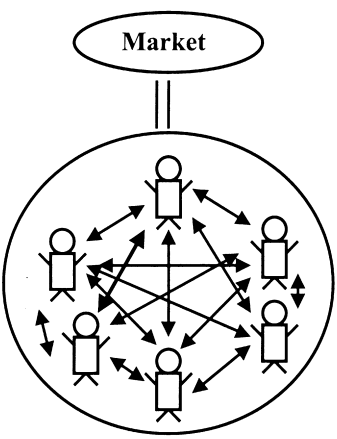

Fig. 1. The market architecture represented by single-population GAs/GP (SGA/SGP).

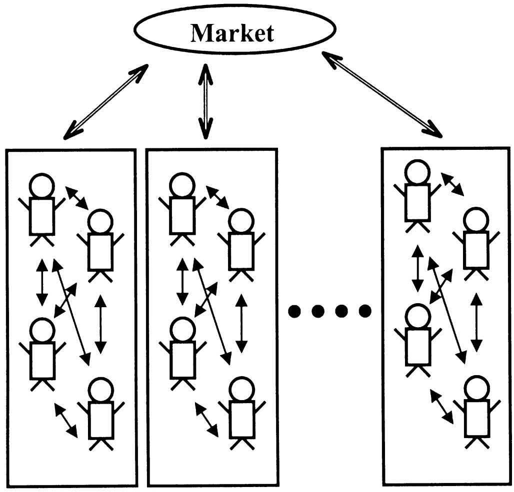

Fig. 2. The market architecture represented by multi-population GAs/GP (MGA/MGP).

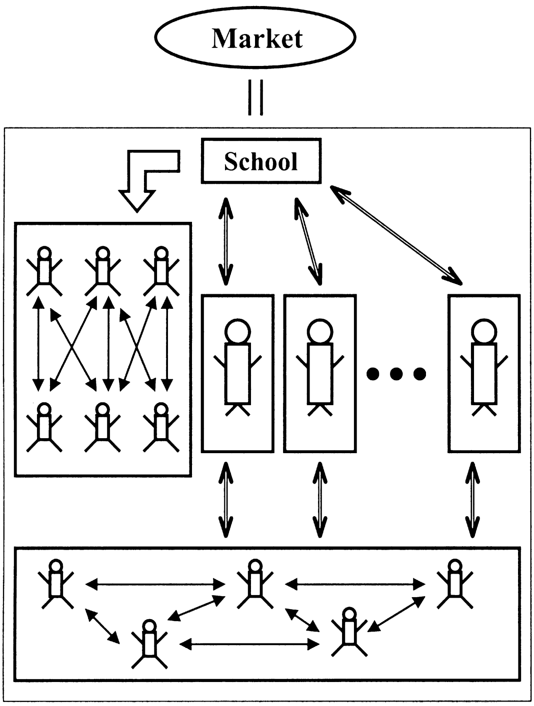

Fig. 3. The market architecture represented by single-population GAs/GP with 'School'.

Fig. 3 represents the architecture of our proposed modification. Again, the market is placed at the top. At the bottom to the right, it is something between Figs. 1 and 2. In the *phenotype*, agents' interaction is direct and identical to what Fig. 1 shows, whereas in the genotype, it looks something like Fig. 2, where there is no direct interaction. The original connection between markets and agents is now replaced by the connection between agents and school shown at the bottom to the left. Inside school, there is again a population of direct interacting agents (faculty), which is pretty much like Fig. 1.

The key elements of our proposed architecture entitled 'MS-GP' (standing for GP implemented with 'School' in the Market) are the procedures '*school*' and the *search*. We shall concretize these procedures with an application to *the artificial stock market*. The artificial stock market is a new but growing field. Some wonders and missions of this research area have been well documented by LeBaron (2000). In his article, LeBaron distinguishes the recent models of complex heterogeneity from those of simple heterogeneity.

> The use of heterogeneous agents is certainly not new to finance, and there is a long history to building heterogeneous agent rational expectations models. What is attempted in this set of computational frameworks is to attack the problem of very complex heterogeneity which leaves the boundary of what can be handled analytically. Traders are made up from a very diverse set of types and behaviors. To make the situation more complex the population of agent types, or the individual behaviors themselves, are allowed to change over time in response to past performance. (p.680)

One of the missions of these agent-based computational models is to *replicate time series features of real markets*. While it will continue to be pursued, the focus of this paper will be much *more fundamental*. As calibration techniques advance, we may expect that sooner or later agent-based financial models will be so powerful that replicating time series features of real markets will not be that daunting. In fact, LeBaron himself has made the following observation:

> Validation remains a critical issue if artificial financial markets are going to prove successful in helping explain the dynamics of real markets. This remains a very weak area for the class of models described here. Further calibration techniques and tighter test will be necessary... . However, there are some key issues which affect these markets in particular. First, they are loaded with parameters *which might be utilized to fit any feature* that is desired in actual data... . (Le Baron, 2000, pp. 698–699, Italics added).

Judging from the results of recent progresses in the literature of artificial stock market, that moment will come in a couple of years. When that moment does come, one may start to question how these calibration techniques can be justified, which leads to the foundation of this research: *can we regard GAs/GP as a suitable model of learning behavior within society?*. The answer can hardly be positive or convincing if Harrald's criticism has not been well taken. We therefore consider this phase of research more fundamental. By saying that, this research tends to modify GP in a manner such that it has a closer connection with human learning and adaptation.

MS-GP brings back *search behavior*, a subject which was once intensively studied in economics but has been largely ignored in the conventional GAs/GP economic literature. As we shall see later, through the idea of *simulated annealing* agents' (traders') search density can be connected to *psychological factors*, such as *peer pressure* or economic factors such as *economic pressure*. Furthermore, the built-in mechanism 'school' enables us to investigate the role of 'school' or the value of 'education' in the evolution of a very specific social process. The statistics generated from simulations, such as the time series of the number of 'students' registered, the number of 'students' who receives futile or fruitful lessons at 'school' can all help us understand how 'school', or information industry in general, coevolves with society.

In Section 2, we shall present the analytical model on which our artificial market is constructed. In Section 3, a concrete application of the institutional GP to the artificial stock market is detailed. Section 4 provides the experimental design. Experiment results and econometric analysis of these designs are given in Section 5 followed by concluding remarks in Section 6.

2. The analytical model

The basic framework of the artificial stock market considered in this paper is the standard asset pricing model (Grossman and Stiglitz, 1980). The market dynamics can be described as an interaction of many heterogeneous agents, each of them, based on her forecast of the future, having the goal to maximize her expected utility. Technically, there are two major components of this market, namely, *traders* and their *interactions*.

2.1. Model of traders

The *trader* part includes *traders*' *objectives* and their *adaptation*. We shall start from traders' motives by introducing their *utility functions*. For simplicity, we assume that all traders share the same utility function. More specifically, this function is assumed to be a *constant absolute risk aversion* (CARA) utility function,

$$U(W_{i,t}) = -\exp(-\lambda W_{i,t}), \tag{1}$$

where $W_{i,t}$ is the wealth of trader i at time period t, and λ is the degree of relative risk aversion. Traders can accumulate their wealth by making investments. There are two assets available for traders to invest. One is the riskless interest-bearing asset called *money*, and the other is the risky asset known as the *stock*. In other words, at each point in time, each trader has two ways to keep her wealth, i.e.,

$$W_{i,t} = M_{i,t} + P_t h_{i,t}, \tag{2}$$

where $M_{i,t}$ and $h_{i,t}$ denotes the money and shares of the stock held by trader i at time t. Given this portfolio $(M_{i,t}, h_{i,t})$, a trader's total wealth $W_{i,t+1}$ is thus

$$W_{i,t+1} = (1+r)M_{i,t} + h_{i,t}(P_{t+1} + D_{t+1}), \tag{3}$$

where P_t is the price of the stock at time period t and D_t is per-share *cash dividends* paid by the companies issuing the stocks. D_t can follow *a stochastic process* not known to traders. Given this wealth dynamics, the goal of each trader is to myopically maximize the one-period expected utility function,

$$E_{i,t}(U(W_{i,t+1})) = E(-\exp(-\lambda W_{i,t+1})|I_{i,t}) \tag{4}$$

subject to

$$W_{i,t+1} = (1+r)M_{i,t} + h_{i,t}(P_{t+1} + D_{t+1}), \tag{5}$$

where $E_{i,t}(.)$ is trader i's conditional expectations of W_{t+1} given her information up to t (the information set $I_{i,t}$), and r is the riskless interest rate.

It is well known that under *CARA* utility and Gaussian distribution for forecasts, trader i's desire demand, $h^*_{i,t+1}$ for holding shares of risky asset is linear in the expected *excess return*:

$$h^*_{i,t} = \frac{E_{i,t}(P_{t+1} + D_{t+1}) - (1+r)P_t}{\lambda \sigma^2_{i,t}}, \tag{6}$$

where $\sigma^2_{i,t}$ is the conditional variance of $(P_{t+1} + D_{t+1})$ given $I_{i,t}$.

One of the essential elements of agent-based artificial stock markets is the formation of $E_{i,t}(.)$, which will be given in detail in the next section.

2.2. Model of price determination

Given $h^*_{i,t}$, the market mechanism is described as follows. Let $b_{i,t}$ be the number of shares trader i would like to submit a bid to buy at period t, and let $o_{i,t}$ be the number trader i would like to offer to sell at period t. It is clear that

$$b_{i,t} = \begin{cases} h^*_{i,t} - h_{i,t-1}, & h^*_{i,t} \geq h_{i,t-1}, \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

and

$$o_{i,t} = \begin{cases} h_{i,t-1} - h^*_{i,t}, & h^*_{i,t} < h_{i,t-1}, \\ 0 & \text{otherwise.} \end{cases} \tag{8}$$

Furthermore, let

$$B_t = \sum_{i=1}^{N} b_{i,t} \tag{9}$$

and

$$O_t = \sum_{i=1}^{N} o_{i,t} \tag{10}$$

be the totals of the bids and offers for the stock at time t, where N is the number of traders. Following Palmer et al. (1994), we use the following simple rationing scheme:[1]

$$h_{i,t} = \begin{cases} h_{i,t-1} + b_{i,t} - o_{i,t} & \text{if } B_t = O_t, \\ h_{i,t-1} + (O_t/B_t)b_{i,t} - o_{i,t} & \text{if } B_t > O_t, \\ h_{i,t-1} + b_{i,t} - (B_t/O_t)o_{i,t} & \text{if } B_t < O_t. \end{cases} \tag{11}$$

All these cases can be subsumed into

$$h_{i,t} = h_{i,t-1} + \frac{V_t}{B_t} b_{i,t} - \frac{V_t}{O_t} o_{i,t}, \tag{12}$$

where $V_t \equiv \min(B_t, O_T)$ is the volume of trade in the stock.

Based on Palmer et al.'s *rationing scheme*, we can have a very simple price adjustment scheme, based solely on the *excess demand* $B_t - O_t$:

$$P_{t+1} = P_t(1 + \beta(B_t - O_t)), \tag{13}$$

where β is a function of the difference between B_t and O_t. β can be interpreted as speed of adjustment of prices. One of the β functions we consider is

$$\beta(B_t - O_t) = \begin{cases} \tanh(\beta_1(B_t - O_t)) & \text{if } B_t \geq O_t, \\ \tanh(\beta_2(B_t - O_t)) & \text{if } B_t < O_t, \end{cases} \tag{14}$$

[1] This simple rationing scheme is chosen mainly to ease the burden of intensive computation. A realistic alternative is to introduce the double auction price mechanism. However, computationally speaking, this idea is very demanding for genetic programming. We are currently working on it in a separate project.

where tanh is the *hyperbolic tangent function*:

$$\tanh(x) \equiv \frac{e^{x} - e^{-x}}{e^{x} + e^{-x}}. \tag{15}$$

Since P_t cannot be negative, we allow the speed of adjustment to be asymmetric to excess demand and excess supply.

The price adjustment process introduced above implicitly assumes that the total number of shares of the stock circulated in the market is fixed, i.e.,

$$H_t = \sum_i h_{i,t} = H. \tag{16}$$

In addition, we assume that dividends and interests are all paid by cash, so

$$M_{t+1} = \sum_i M_{i,t+1} = M_t(1 + r) + H_t D_{t+1}. \tag{17}$$

2.3. Model of adaptive traders

In this section, we shall address the formation of traders' expectations, $E_{i,t}(P_{t+1} + D_{t+1})$ and $\sigma^2_{i,t}$. Motivated by the *martingale hypothesis* in finance, we shall assume the following function form for $E_{i,t}(\,.\,)$:

$$E_{i,t}(P_{t+1} + D_{t+1}) = (P_t + D_t)(1 + \theta_1 \tanh(\theta_2 \cdot f_{i,t})). \tag{18}$$

The virtue of this function form is that, if $f_{i,t} = 0$, then the trader actually validates the martingale hypothesis. Therefore, from the *cardinality* of set $\{i \,|\, f_{i,t} = 0\}$, denoted by $N_{1,t}$, we can know *how well the efficient market hypothesis is accepted among traders*.[2] The population of functions $f_{i,t}$ $(i = 1, \dots, N)$ is determined by the genetic programming procedure *Business School* and *Search in Business School* given in the following two subsections.

As to the subjective risk equation, we take a modification of the equation originally used by Arthur et al. (1997).

$$\sigma^2_{i,t} = (1 - \theta_3)\sigma^2_{t-1|n_1} + \theta_3[(P_t + D_t - E_{i,t-1}(P_t + D_t))^2], \tag{19}$$

[2] Of course, this number cannot be taken without some caution. Notice that $f_{i,t}$ can be arbitrarily close to zero, but not identically zero. In this case, it is *essentially* consistent with the martingale hypothesis, and $N_{1,t}$ might not be completely informative. To cope with this problem, what our programming does is to treat $f_{i,t}$ *essentially zero* if

$$|f_{i,t} - 0| \leq \vartheta$$

and ϑ can be determined by the end-user. In this paper, ϑ is set to 10^{-5}.

where

$$\sigma^2_{t|n_1} = \frac{\sum_{j=0}^{n_1-1}[P_{t-j} - \bar{P}_{t|n_1}]^2}{n_1 - 1} \tag{20}$$

and

$$\bar{P}_{t|n_1} = \frac{\sum_{j=0}^{n_1-1} P_{t-j}}{n_1}. \tag{21}$$

In other words, $\sigma^2_{t-1|n_1}$ is simply the *historical volatility* based on the past n_1 observations.

2.4. Business school and single-population GP

The major component of artificial stock markets is the adaptive traders, who can be regarded as an evolving population. Since Arifovic (1994), genetic algorithm has been employed to drive the evolving population of agents in economics. Chen and Yeh (1996) generalized this approach by using genetic programming. The style of GP used in Chen and Yeh (1996) is known as *single-population* GP in *agent-based computational economics*, which is different from *multi-population GP*. In single-population GP, each tree can be regarded as a forecasting model used by an agent; hence, the adaptation of agents (in terms of their forecasting models) can be directly represented by the standard operation of GP. However, due to Harrald's criticism mentioned in Section 1, we consider a modified version of single-population GP in this paper. Our modified version is characterized as an addition of a *business school* to the artificial stock market.

The business school in our model functions as usual business schools in the real world. It mainly consists of faculty, and their different kinds of models (schools of thoughts). Let F be the number of faculty members (forecasting models). These models are propagated via a competition process driven by the faculty through publications. In this *academic* world, a scholar can ill afford to keep something serious to herself if she wants to be well acknowledged. If we consider business school a collection of forecasting models, then we may well use single-population GP to model its adaptation.

Nonetheless, scholars and traders may care about different things. Therefore, in this paper, different fitness functions are chosen to take care of such a distinction. For scholars, the fitness function is chosen purely from a scientific viewpoint, say, forecasting accuracy. For example, one may choose *mean absolute percentage error* (MAPE) as the fitness function (Table 1). Single-population GP is then conducted in a standard way. Each faculty member (forecasting model) is represented by a tree (GP parse tree). The faculty will be evaluated with

Table 1
Parameters of the stock market[a]

The stock market	
Shares of the stock (H) per capita	1
Initial money supply (M_1) per capita	100
Interest rate (r)	0.1
Stochastic process (D_t)	IID $\sim$ *Uniform*(5.01, 14.99)
Price adjustment function	tanh
Price adjustment (β_1)	10^{-4}
Price adjustment (β_2)	0.2×10^{-4}
Business school	
Number of faculty members (F)	500
Number of trees created by the full method	50
Number of trees created by the grow method	50
Function set	$\{+, -, *, /, \text{Sin}, \text{Cos}, \text{Exp}, \text{R log}, \text{Abs}, \text{Sqrt}\}$
Terminal set	$\{P_t, P_{t-1}, \ldots, P_{t-10}, P_{t-1} + D_{t-1}, \ldots, P_{t-10} + D_{t-10}\}$
Selection scheme	Tournament selection
Tournament size	2
Probability of creating a tree by reproduction (p_r)	0.10
Probability of creating a tree by crossover (p_c)	0.70
Probability of creating a tree by mutation (p_m)	0.20
Probability of mutation	0.0033
Probability of leaf selection under crossover	0.5
Mutation scheme	Tree mutation
Replacement scheme	Tournament selection
Maximum depth of tree	17
Number of generations	20,000
Maximum number in the domain of Exp	1700
Criterion of fitness (Faculty members)	MAPE
Evaluation cycle (m_1)	20
Sample Size (MAPE) (m_2)	10
Traders	
Number of Traders (N)	500
Degree of RRA (λ)	0.5
Criterion of fitness (Traders)	Increments in wealth (Income)
Sample size of $\sigma^2_{t\|n_1}$ (n_1)	10
Evaluation cycle (n_2)	1
Sample size (n_3)	10
Search intensity (I^*)	5
θ_1	0.5
θ_2	10^{-4}
θ_3	0.0133

[a]The number of trees created by the full method or grow method is the number of trees initialized in Generation 0 with the depth of tree being 2–6. For details, see Koza (1992).

a prespecified schedule, say once for every m_1 trading days. The review procedure proceeds as follows.

At the evaluation date, say t, each forecasting model (faculty member) will be reviewed by a *visitor*. The visitor is another model which is generated randomly from the collection of the existing models in the business school at $t-1$, denoted by GP_{t-1}, by one of the following three genetic operators, reproduction, crossover and mutation, each with probability p_r, p_c, and p_m (Table 1). In the case of reproduction or mutation, we first randomly select two GP trees, say, $gp_{j,t-1}$ and $gp_{k,t-1}$. The MAPE of these two trees over the last m_2 days' forecasts are calculated. A *tournament selection* is then applied to these two trees. The one with lower MAPE, say $gp_{j,t-1}$, is selected. We then run a tournament again over the host $gp_{i,t-1}$ and the visitor $gp_{j,t-1}$ (in the case of reproduction) or $gp'_{j,t-1}$ (in the case of mutation) based on the criterion MAPE, and $gp_{i,t}$ is the winner of this tournament.

In the case of crossover, we first randomly select two pairs of trees, say $(gp_{j_1,t-1}, gp_{j_2,t-1})$ and $(gp_{k_1,t-1}, gp_{k_2,t-1})$. The tournament selection is applied separately to each pair, and the winners are chosen to be parents. The children, say (gp_1, gp_2), are born. One of them is randomly selected to compete with $gp_{i,t-1}$, and the winner is $gp_{i,t}$. The following is a pseudo-program of the procedure *Business School* (also see Fig. 4). Table 1 is an example of the specification of the control parameters to evolve the business school.

Procedure [Business School]

0. **begin**
   1. Calculate $MAPE(gp_{i,t-1})$
   2. A = **Random**(R, C, M) with (p_r, p_c, p_m)
   3. If A = C, go to step (11).
   4. $(gp_1, gp_2) =$ (**Random**(GP_{t-1}), **Random**(GP_{t-1}))
   5. Calculate $MAPE(gp_1)$ and $MAPE(gp_2)$.
   6. $gp_{new} =$ **Tournament Selection** $(MAPE(gp_1), MAPE(gp_2))$
   7. If A = R, go to step (17).
   8. $gp_{new} \leftarrow$ **Mutation**(gp_{new})
   9. Calculate $MAPE(gp_{new})$
   10. Go to step (17)
   11. Randomly select two pairs of trees from GP_{t-1}
   12. Calculate MAPE of these two pairs of GP trees
   13. $gp_1 =$ **Tournament Selection** (PAIR 1)
   14. $gp_2 =$ **Tournament Selection** (PAIR 2)
   15. $(gp_1, gp_2) \leftarrow$ **Crossover**(gp_1, gp_2)
   16. $gp_{new} =$ **Random**(gp_1, gp_2)
   17. $gp_{i,t} =$ **Tournament Selection** $(MAPE(gp_{i,t-1}), MAPE(gp_{new}))$
18. **end**

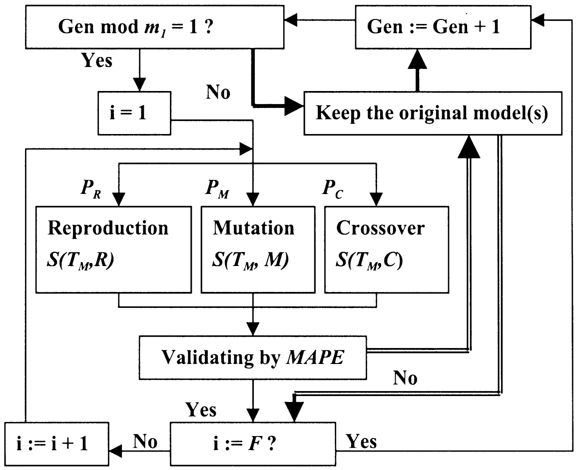

T_M: Tournament selection according to MAPE
$S(T_M,i)$: Selection procedure according to tournament selection with the criterion MAPE based on i (genetic operator)
i : R, M, C represent reproduction, mutation and crossover respectively
m_I : Evaluation cycle
F : Number of faculty member

Fig. 4. Evolution of the Business School.

2.5. Traders and business school

Given the adaptive process of the business school, the adaptive process of traders can be described as a sequence of two decisions. First, should she go back to the business school to *take classes*? Second, should she *follow* the lessons learned at school? In the real world, the first decision somehow can be more *psychological* and has something to do with *peer pressure*. One way to model the influence of peer pressure is to suppose that each trader will examine how well she has performed over the last n_2 trading days, when compared with other traders. Suppose that traders are ranked by *the net change of wealth* over the last n_2 trading days. Let $W_{i,t}^{n_2}$ be this net change of wealth of trader i at time period t, i.e.,

$$\Delta W_{i,t}^{n_2} \equiv W_{i,t} - W_{i,t-n_2} \tag{22}$$

and, let $R_{i,t}$ be her rank. Then, the probability that trader i will go to business school at the end of period t is assumed to be determined by

$$p_{i,t} = \frac{R_{i,t}}{N}. \tag{23}$$

The choice of the function $p_{i,t}$ is quite intuitive. It simply means that

$$p_{i,t} < p_{j,t} \quad \text{if } R_{i,t} < R_{j,t}. \tag{24}$$

In words, the traders who come out top shall suffer less peer pressure, and hence have less motivation to go back to school than those who are ranked at the bottom.

In addition to peer pressure, a trader may also decide to go back to school out of a sense of *self-realization*. Let the growth rate of wealth over the last n_2 days be

$$\delta_{i,t}^{n_2} = \frac{W_{i,t} - W_{i,t-n_2}}{|W_{i,t-n_2}|} \tag{25}$$

and let $q_{i,t}$ be the probability that trader i will go back to business school at the end of the tth trading day, then it is assumed that

$$q_{i,t} = \frac{1}{1 + \mathrm{e}^{\delta_{i,t}^{n_2}}}. \tag{26}$$

The choice of this density function is also straightforward. Notice that

$$\lim_{\delta_{i,t}^{n_2} \to \infty} q_{i,t} = 0 \tag{27}$$

and

$$\lim_{\delta_{i,t}^{n_2} \to -\infty} q_{i,t} = 1. \tag{28}$$

Therefore, the traders who have made great progress will naturally be more confident and hence have little need for schooling, whereas those who suffer devastating regression will have a strong desire for schooling.

In sum, for trader i, the decision to go to school can be considered as a result of *a two-stage independent Bernoulli experiments*. The success probability of the first experiment is $p_{i,t}$. If the outcome of the first experiment is success, the trader will go to school. If, however, the outcome of the first experiment is failure, the trader will continue to carry out the second experiment with the success probability $q_{i,t}$. If the outcome of the second experiment is success, then the trader will also go to school. Otherwise, the trader will quit school. If we let $r_{i,t}$ be the probability that trader i decides to go to school, then

$$\begin{aligned} r_{i,t} &= p_{i,t} + (1 - p_{i,t})q_{i,t} \\ &= \frac{R_{i,t}}{N} + \frac{N - R_{i,t}}{N} \frac{1}{1 + \mathrm{e}^{\delta_{i,t}^{n_2}}}. \end{aligned} \tag{29}$$

Once a trader decides to go to school, she has to make a decision on what kinds of classes to take. Since we assume that business school, at period t, consists of 500 faculty members (forecasting models), let us denote them by $gp_{j,t}$ ($j = 1, 2, \ldots, 500$). The class-taking behavior of traders is assumed to follow the following sequential search process. The trader will randomly select one *forecasting model* $gp_{j,t}$ ($j = 1, \ldots, 500$) with a uniform distribution. She will then *validate* this model by using it to fit the stock price and dividends over the last n_3 trading days, and compare the result (MAPE) with her original model. If it outperforms the old model, she will discard the old model, and put the new one into practice. Otherwise, she will start another random selection, and do it again and again until either she has a successful search or she continuously fail I^* times. The following is a pseudo program of the procedure *Visiting the Business School* (also see Fig. 5).

Procedure [Visiting Business School]

0. **begin**
   1. Calculate $MAPE(f_{i,t})$
   2. $I \leftarrow 1$
   3. Randomly select a $gp_{j,t}$ ($\sim U[1, 500]$)
   4. Calculate $MAPE(gp_{j,t})$
   5. If $MAPE(gp_{j,t}) < MAPE(f_{i,t})$, go to Step (10)
   6. $I \leftarrow I + 1$
   7. If $I \leq I^*$, go to step (3)
   8. $f_{i,t+1} = f_{i,t}$
   9. Go to Step (11)
   10. $f_{i,t+1} = gp_{j,t}$
11. **end**

Eq. (29) and the procedure *Visiting Business School* give the distinguishing feature of our adaptive traders. As we mentioned earlier, there is no direct interaction among traders in terms of the genotype. Therefore, the conventional SGA or SGP used to evolve a population of traders is no longer applicable here. In other words, our traders are not GP(GA)-based. Instead, their adaptation behavior is modeled by an explicit search process. The search process starts with *a decision to search or not*. This decision is *stochastic*, i.e., the trader at any point in time cannot be sure whether she should start searching, and the uncertainty of this decision is further modeled by a technique similar to *simulated annealing* (SA).[3] In sum, it is a society composing of SA-based traders and SGP-based faculty, who coevlove with different fitness functions (objective functions).

[3] To be precise, the search procedure introduced above is not simulated annealing. In simulated annealing, the decision to accept a new solution is random. However, here, it is *the decision to search* a random one. By this setting, the learning rate is endogenously determined for each individual rather than exogenously given.

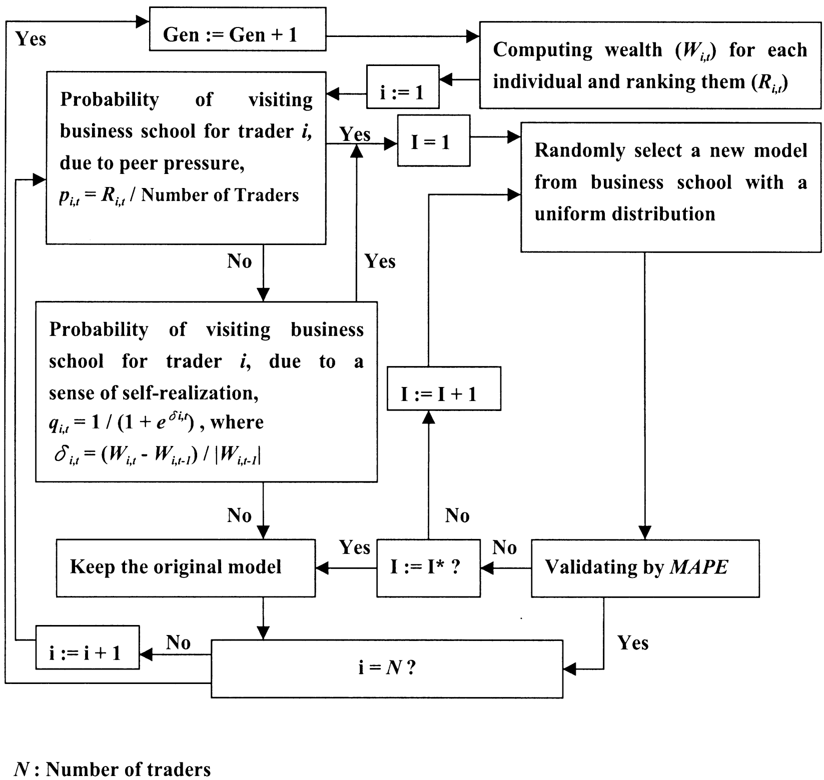

Fig. 5. Trader's search process in the Business School.

3. Experimental designs

One of the formidable tasks for agent-based computational stock markets is the *design of traders*. As LeBaron (2000) pointed out: 'The computational realm has the advantages and disadvantages of a wide open space in which to design traders, and new researchers should be aware of the daunting design questions that they will face. Most of these questions still remain relatively unexploited at this time'. (p. 696) Nevertheless, one should notice that this issue is not confined to agent-based computational finance, and is widely shared by all research in *bounded rationality*. For example, Sargent (1993) stated 'This area is wilderness because the research faces so many choices after he decides to forgo the discipline provided by equilibrium theorizing'. (p. 2)

LeBaron's and Sargent's description of this wilderness can be further exemplified by Table 1. Facing such *a wide open space*, we have to admit that some choices we made may be arbitrary, and that the results may not be robust to all designs. Hence, in addition to run many runs in a single design, it is also crucial to test many different designs, i.e., to test many tables like Table 1, while with different parameters. This paper, however, has a very limited scope, i.e., to illustrate the rich dynamics our artificial stock market can possibly offer, and the questions it can effectively deal with a single experiment.[4] Therefore, while we do fine-tune some of our parameters listed in Table 1, we do not intentionally *calibrate* our parameters for the purpose of replicating the *stylized facts* of financial time series.[5]

The simulation results of our artificial stock market are mainly a series of time series variables of *traders* (*microstructure*) and the *market*. They are summarized in Table 2.

4. Simulation results

Based on the experimental design given above (Table 1), a single run with 14,000 generations was conducted. Notice that the number of generations is also the time scale of simulation, i.e., $GEN = t$. In other words, we are simultaneously evolving the population of traders while deriving the price P_t. In the following, we shall present our results in an order to answer a series of questions motivated by Pagan (1996).[6]

1. Are prices and returns normally distributed?
2. Does the price series have a *unit root*, i.e., does the price series follow an $I(1)$ process?
3. Are returns independently and identically distributed?

In addition to the '*up*' part, agent-based computational models provide us with rich opportunity to study the microstructure, i.e., the behavioral aspect of traders. In our artificial stock market, a trader's behavior can be well kept track

[4] In fact, we did have multiple runs before coming to this point. The simulations results presented here can be considered a typical one in the sense that many properties reported here are widely observed in other runs. Therefore, to make the presentation to have a focus, we shall only exemplify what we learn from this market by a single 'typical' run. Nevertheless, to have better communication with readers, we have put our program to the website so that it can be freely downloaded, examined and tested. The website address is:

http://econo.nccu.edu.tw/ai/staff/csh/Software.htm

[5] According to LeBaron (2000), one of the missions of the agent-based modeling of financial markets is to replicate time series features of real markets. Lux (1995, 1998), Lux and Marchesi (1999), Chen and Kuo (1999), Chen et al. (2000b) have showed how these stylized facts can be replicated in a specific style of agent-based models.

[6] Pagan (1996) summarized a list of *stylized facts* in financial time series.

Table 2
Time series generated from the artificial stock market

Aggregate variables	
Stock price	P_t
Trading volumes	V_t
Totals of the bids	B_t
Totals of the offers	O_t
No. of martingale believers	$N_{1,t}$
No. of traders registered to Business School	$N_{2,t}$
No. of traders with successful search in Business School	$N_{3,t}$
Individual trader	
Forecasts	$f_{i,t}$
Subjective risks	$\sigma_{i,t}$
Bid to buy	$b_{i,t}$
Offer to sell	$o_{i,t}$
Wealth	$W_{i,t}$
Income	$\Delta W^1_{i,t}$
Rank of profit-earning performance	$R_{i,t}$
Complexity (depth of $f_{i,t}$)	$k_{i,t}$
Complexity (No. of nodes of $f_{i,t}$)	$\kappa_{i,t}$

of by a list of variables given in Table 2. This list of variables enables us to address a lot of interesting issues in behavioral finance.

1. What does the traders actually believe? Does she believe in the efficient market hypothesis?
2. What exactly is the forecasting model (or the trading strategy) employed by the trader?
3. How sophisticated is the trader? Will she get more and more sophisticated as time goes on?

In the following, we shall illustrate how these issues can be approached by our agent-based artificial stock market.

First, are price and returns normally distributed? The time series plot of the stock price is drawn in Fig. 6. Over this long horizon, P_t fluctuates between 55 and 105. The basic statistics of this series, $\{P_t\}_{t=1}^{14\,000}$, is summarized in Table 3. Given the price series, the return series is derived as usual

$$r_t = \ln(P_t) - \ln(P_{t-1}). \tag{30}$$

Fig. 7 is a time series of stock return, and Table 4 gives the basic statistics of this return series. From these two tables, neither the stock price series $\{P_t\}$ nor return series $\{r_t\}$ is normal. The null hypothesis that these series are normal are rejected by the Jarqu–Bera statistics in all periods. The fat-tail property is

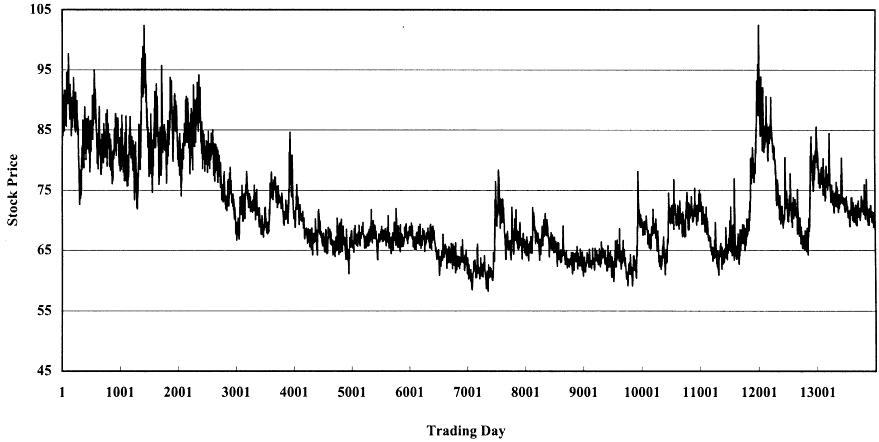

Fig. 6. Time series plot of the stock price.

Table 3
Basic statistics of the artificial stock price series

Periods	$\bar{P}$	σ	Skewness	Kurtosis	Jarqu-Bera	p-value
1–2000	84.07	4.82	0.34	3.07	40.62	0.00
2001–4000	76.43	5.84	0.65	2.60	153.49	0.00
4001–6000	67.28	1.84	0.94	5.07	654.75	0.00
6001–8000	65.17	3.27	0.67	3.85	212.46	0.00
8001–10000	64.46	2.49	1.16	5.28	887.91	0.00
10001–12000	68.44	5.09	2.24	11.46	7660.11	0.00
12001–14000	74.57	5.48	1.00	3.71	381.93	0.00

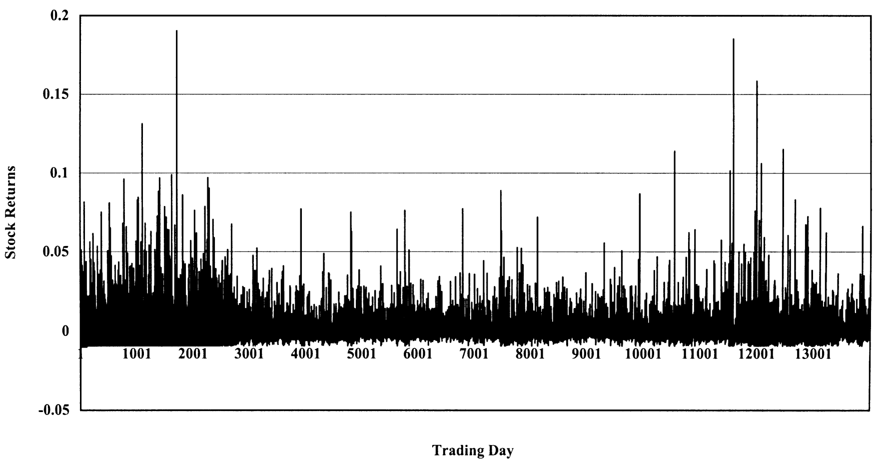

Fig. 7. Time series plot of stock returns.

Table 4
Basic statistics of the artificial stock return series

Periods	$\bar{P}$	σ	Skewness	Kurtosis	Jarqu-Bera	p-value
1–2000	− 0.000074	0.015	3.53	23.64	39676.46	0.00
2001–4000	− 0.000057	0.010	3.26	18.83	24461.55	0.00
4001–6000	− 0.000018	0.007	3.72	25.94	48486.08	0.00
6001–8000	− 0.000024	0.007	3.70	25.79	47869.55	0.00
8001–10000	0.000032	0.007	3.69	26.97	52452.04	0.00
10001–12000	0.000169	0.010	6.91	86.56	597871.50	0.00
12001–14000	− 0.000154	0.009	4.18	32.80	79867.54	0.00

Table 5
Unit root test and PSC filtering[a]

Periods	DF of P_t	(p, q)
1–2000	− 0.285	(0, 0)
2001–4000	− 0.288	(0, 0)
4001–6000	− 0.150	(0, 0)
6001–8000	− 0.180	(0, 0)
8001–10000	0.173	(0, 0)
10001–12000	0.680	(0, 0)
12001–14000	− 0.753	(0, 0)

[a] The MacKinnon critical values for rejection of hypothesis of a unit root at 1% (5%) significance level is − 2.5668 (− 1.9395).

especially striking in the return series. This result is consistent with one of stylized facts documented in Pagan (1996).

Second, does the price series have a unit root? The standard tool to test for the presence of a unit root is the celebrated Dickey–Fuller (DF) test (Dickey and Fuller, 1981). The DF test consists of running a regression of the first difference of the log prices series against the series lagged once.

$$\Delta \ln(P_t) = \ln(P_t) - \ln(P_{t-1}) = \beta_1 \ln(P_{t-1}). \tag{31}$$

The null hypothesis is that β_1 is zero, i.e., $\ln(P_t)$ contains a unit root. If β_1 is significantly different from zero then the null hypothesis is rejected. As can be seen from the second column of Table 5, from the total number of 7 periods none leads to a rejection of the presence of a unit root.

Third, are returns independently and identically distributed? Here, we followed the procedure of Chen et al. (2000b). This procedure is composed of two steps, namely, *the PSC filtering* and *the BDS testing*. We first applied the Rissanen's predictive stochastic complexity (**PSC**) to filter the linear process. The third column of Table 5 gives us the $ARMA(p, q)$ process extracted from the return series. Interestingly enough, all these seven periods are *linearly independent* ($p = 0$, $q = 0$). This result corresponds to the classical version of *the efficient market hypothesis*.

Once the linear signals are filtered, any signals left in the residual series must be *nonlinear*. Therefore, one of the most frequently used statistic, the BDS test, is applied to the residuals from the PSC filter. Since none of the seven return series have linear signals, the BDS test is directly applied to the original return series. There are two parameters required to conduct the BDS test. One is the distance parameter (ε standard deviations), and the other is the *embedding dimension* (*DIM*). We found the result is not sensitive to the first choice, and hence, we only report the result with $\varepsilon = 1$. As to the embedding dimension, we tried $DIM = 2, 3, 4, 5$, and the result is given in Table 6. Since the BDS test is asymptotically normal, it is quite easy to have an eyeball check on the results.

What is a little surprising is that the null hypothesis of IID (identically and independently distributed) is not rejected in 6 out of 7 periods. The only period whose return series has nonlinear signals is Period 5. Putting the result of PSC filtering and BDS testing together, our return series is *efficient* to the degree that, 85% of the time, it can be regarded as a iid series. But, if the series is indeed independent (no signals at all), *what is the incentive for traders to search*? Clearly, here, we have come to the issues raised by Grossman and Stiglitz 20 years ago (Grossman and Stiglitz, 1980).

One of the advantages agent-based computational economics (the bottom-up approach) is that it allows us to observe *what traders are actually thinking and doing*. Are they *martingale believers*? That is, do they believe that

$$E_t(P_{t+1} + D_{t+1}) = P_t + D_t?. \tag{32}$$

If they do not believe in the martingale hypothesis, do they search intensively? In other words, do they go to school and can still learn something useful in such an

Table 6
BDS test[a]

Periods	DIM = 2	DIM = 3	DIM = 4	DIM = 5	Reject
1–2000	− 0.36	− 0.20	− 0.14	− 0.18	No
2001–4000	− 0.16	0.13	0.40	0.57	No
4001–6000	1.34	1.35	1.22	1.24	No
6001–8000	0.89	0.99	1.18	1.35	No
8001–10000	1.93	2.38	2.64	2.69	Yes
10001–12000	0.85	0.92	0.96	0.87	No
12001–14000	0.29	0.21	0.37	0.66	No

[a]The test statistic is asymptotically normal with mean 0 and SD 1. The significance level of the test is set at 0.05.

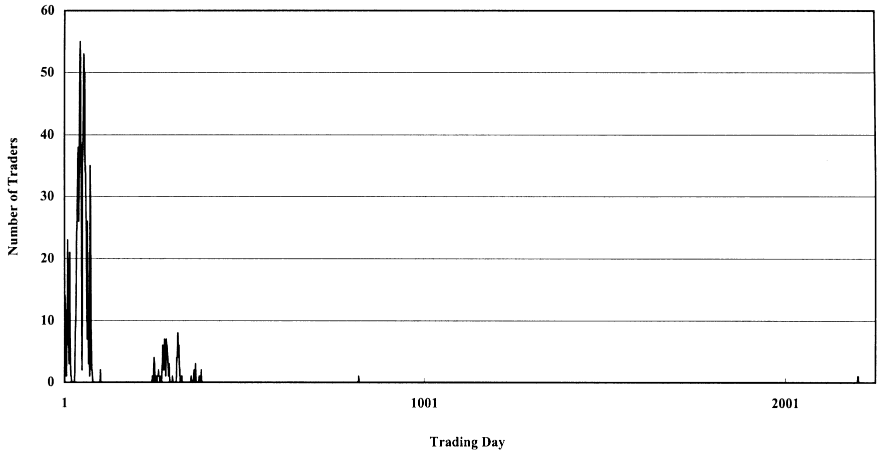

Fig. 8. The number of traders with martingale strategies on each trading day.

iid-series environment? To answer the first question, the time series of $N_{1,t}$ in Table 2 is drawn in Fig. 8. The figure is drawn only up to the first 2203 trading days, because after that the group of believers goes extinct. Hence, while econometricians may claim that the return series is iid, traders simply do not buy it.[7]

This naturally brings up the second question: *if they do not believe in the martingale hypothesis*, *what do they actually do*? Fig. 9 is the time series plot of the number of traders with successful search, $N_{3,t}$. Due to the density of the plot and the wide range of fluctuation, this figure is somewhat complicated and difficult to read. We, therefore, report the average of $N_{3,t}$ over different periods of trading days in Table 7. From Table 7, it can be seen that the number of traders with successful search, on the average, fluctuates about 200. At a rough estimate, 40% of the traders benefit from business school per trading day. Clearly, search in business school is not futile.

It is interesting to know what kind of useful lessons traders learn from business school. Based on our design given in Section 3, what business school offers is a collection of forecasting models $\{gp_{i,t}\}$, which may well capture *the recent movement of the stock price and dividends*. Therefore, while *in the long-run* the return series is iid, traders under survival pressures do not care much about

[7] One possible explanation for this inconsistency is that, under survival pressure, traders only care about their *short-term performance*, and are only looking for models which are able to work in the short term. In fact, as what we shall see later, traders in our artificial world are very active on trying and using new models.

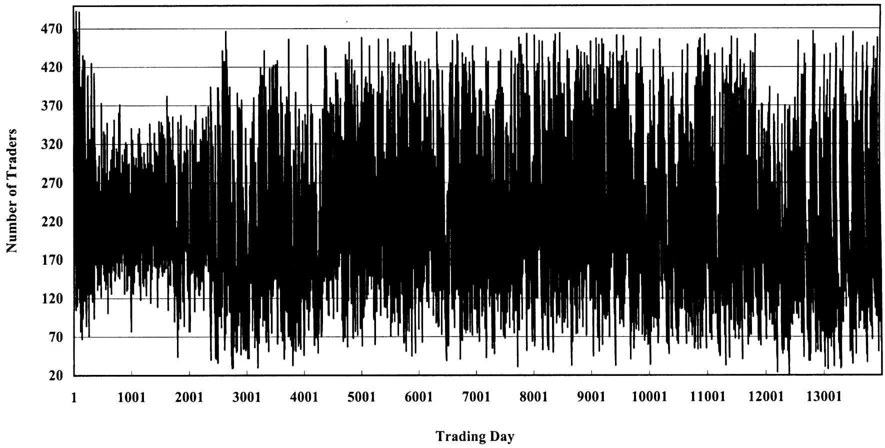

Fig. 9. The number of traders with successful search on each trading day.

Table 7
Microstructure statistics: average of traders with successful search and complexity of evolving strategies[a]

Periods	$\bar{N}_3$	$\bar{k}$	$\bar{\kappa}$
1–2000	209.13	17.85	8.14
2001–4000	189.03	28.14	9.66
4001–6000	218.53	54.34	13.29
6001–8000	215.91	59.51	14.13
8001–10000	220.78	76.60	14.74
10001–12000	206.80	69.22	13.97
12001–14000	185.40	50.58	12.94

[a] $\bar{N}_3$ is the average of $N_{3,t}$ taken over each period. $\bar{k}$ and $\bar{\kappa}$ are the average of k_t and κ_t taken over each period.

this long-run property. What motivates them to search and helps them to survive is in effect *brief signals*. A similar observation was made by Peters (1991):

> The evidence calls into question the Efficient Market Hypothesis, which underlies the linear mathematics used in most capital market theory. It also lends validity to a number of investment strategies that *should not work* if markets are efficient, … . This finding is of particular importance for practitioners, because experience has shown that these strategies *do work* when *properly* applied, even though theory tells us *they should not work in a random-walk environment.* (italics added)

Another way to see what traders may learn from business school is to examine the forecasting models they employ. However, this is a very large database, and is difficult to deal with directly. But, since all forecasting models are in the format of LISP trees, we can at least ask *how complex these forecasting models are*. To do so, we give two definitions of the *complexity* of a GP-tree. The first definition is based on the *number of nodes* appearing in the tree, while the second is based on the *depth* of the tree. On each trading day, we have a profile of the evolved GP-trees for 500 traders, $\{f_{i,t}\}$. The complexity of each tree is computed. Let $k_{i,t}$ be *the number of nodes* of the model $f_{i,t}$ and $\kappa_{i,t}$ be the *depth* of $f_{i,t}$. We then average as follows

$$k_t = \frac{\sum_i^{500} k_{i,t}}{500} \quad \text{and} \quad \kappa_t = \frac{\sum_i^{500} \kappa_{i,t}}{500}. \tag{33}$$

Figs. 10 and 11 are the time series plots of k_t and κ_t. One interesting hypothesis one may make is that *the degree of traders' sophistication is an increasing function of time* (*the monotone hypothesis*). In other words, traders will evolve to be more and more sophisticated as time goes on. However, this is not the case here. Both figures evidence that, while traders can evolve toward a higher degree of sophistication, at some point in time, they can be simple as well (also see Table 7). Despite the rejection of the monotone hypothesis, we see no evidence that traders' behavior will converge to the simple martingale model.

Figs. 9–11 together leave us an impression that traders in our artificial stock market are very adaptive. About this phenomenon, Arthur (1992) conducted a survival test on it.

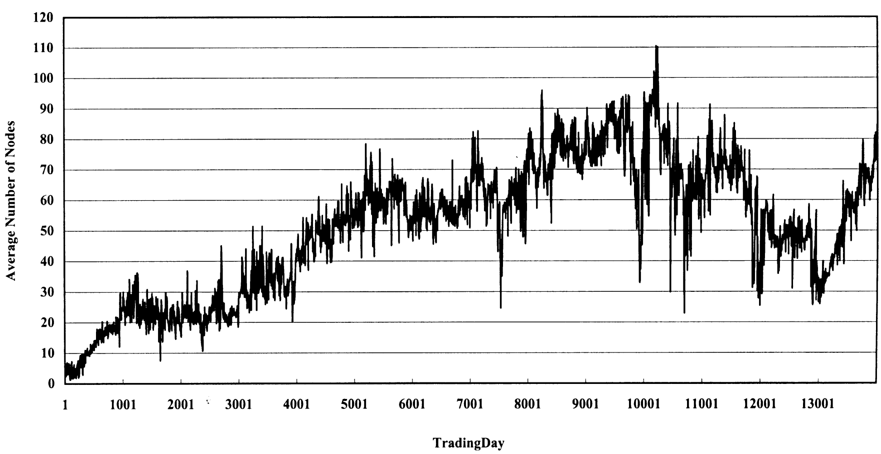

Fig. 10. Trading's complexity: the average of the number of nodes of GP-trees.

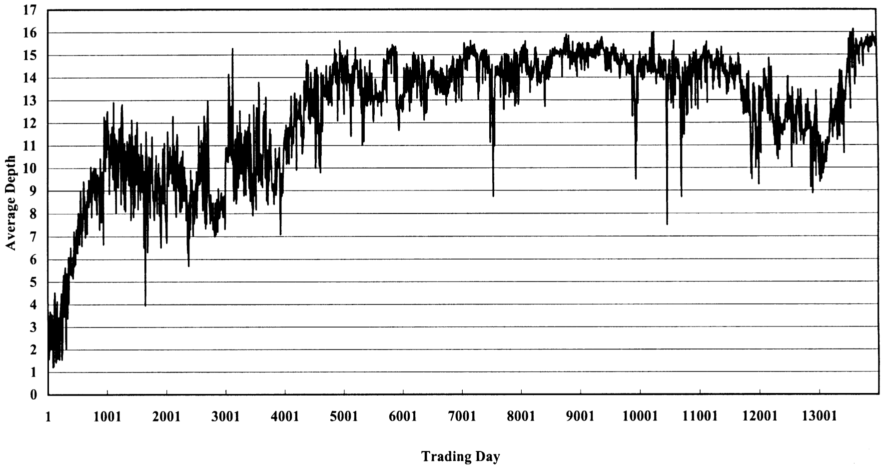

Fig. 11. Trader's complexity: the average of depth of GP-trees.

> We find no evidence that market behavior ever settles down; the population of predictors continually co-evolves. One way to test this is to take agents out of the system and inject them in again later on. *If market behavior is stationary they should be able to do as well in the future as they are doing today.* But we find that when we 'freeze' a successful agent's predictors early on and inject the agent into the system much later, the formerly successful agent is now a dinosaur. His predictions are *unadapted* and perform poorly. *The system has changed.* (p. 24, italics added)

Arthur's interesting experiment can be considered as a measure of the speed of change in a system. If a system changes in a very fast manner, then knowledge about the system has to be updated in a similar pace; otherwise, the knowledge acquired shall soon become obsolete. To see how fast our artificial stock market changes, we made an experiment *similar to* Arthur's survival test. Since our artificial market business school updates every 20 periods ($m_1 = 20$, Table 1), we can measure how fast the knowledge become obsolete by calculating the number of traders with successful search on the hth day after business school has updated the knowledge.

It is expected that knowledge acquired on the day immediately after the updating day should be most helpful for the searching traders. Therefore, the number of traders with successful search should be strikingly high on that day, and the farther it is from the updating, the less the chance of having a successful search. More precisely, denote $N_{3,t}$ by N_{3,h_i}, where $t = (i)\cdot 20 + h$, and let

$$N_{3,h} = \frac{\sum_{i=1}^{14\,000/20} N_{3,h_i}}{14\,000/20} \tag{34}$$

Table 8
Average of the number of traders with successful search on the h day after business school has updated the information

h	$\bar{N}_{3,h}$	h	$\bar{N}_{3,h}$	h	$\bar{N}_{3,h}$	h	$\bar{N}_{3,h}$
1	308.52	6	208.88	11	189.87	16	183.49
2	270.24	7	200.80	12	188.04	17	184.49
3	246.39	8	196.56	13	187.81	18	186.54
4	230.82	9	193.27	14	187.94	19	193.39
5	218.86	10	191.47	15	184.61	20	185.39

then a test *similar to* Arthur's 'Jurassic Park' experiment can be reformulated as follows. $N_{3,h}$ *is a monotonic decreasing function of h.* To see whether this property will apply to our system, Table 8 reports the statistics $N_{3,h}$. This series of numbers starts with a peak at 308, and quickly goes down below 300 and then drops further below 200 as h increases. This result simply says that when more and more people knows the secret, there can be no longer any secret.

The last result also shows the co-evolving complex dynamics between business school and the market. To survive, school must adapt to market dynamics. On the other hand, market dynamics generate *students* for 'school' who, in turn, bring the knowledge learned from 'school' back to the market, and that knowledge may have further impact on market dynamics. The patterns discovered by business school are eventually annihilated by the traders who learn and make a living on these patterns. However, on the process of annihilating these patterns, new patterns are further generated for school to discover, and this process goes on and on. One may call this process a *self-destruction and generation process*.

5. Concluding remarks

The single experiment conducted here has demonstrated the rich dynamics that our artificial stock market can generate. We also show the relevance of this rich dynamics to *financial econometrics* and *behavioral finance*. For the latter, we address Peters' criticism on the efficient market hypothesis as well as the survival test with our dynamics of microstructure. It is interesting to note that, while econometricians on the *top* may claim that our artificial market is efficient, our traders on the *bottom* do not act as if they believe in the efficient market hypothesis. This result seems to be consistent with our experience of the real world, and is one of the interesting features one may expect from the bottom-up approach.

Appendix. The software AIE-ASM

The software used to simulate the artificial stock market reported here can be directly downloaded from http://econo.nccu.edu.tw/ai/staff/csh/Software.htm.

The current version is *AIE-ASM* Version 2. This website will be continuously updated when new versions are available. There are two papers which can be helpful for potential users. Chen et al. (2000a) provided instructions on how to install this software. They further exemplified the use of this software system by examining its performance from the aspect of *price discovery*. Chen and Yeh (2000) studied the complex dynamics of this software system as a 25-dimensional dynamical system under several different settings, which include the ones with and without a b-school, ones with different speed of price adjustment (β_1 and β_2), ones with different rank functions (Eq. (23)). Apart from research, for those who are interested in using this software system for teaching, there are also materials written in Power Point, which are available upon request.

References

Andrews, M., Prager, R., 1994. Genetic programming for the acquisition of double auction market strategies. In: Kinnear, K.E. (Ed.), Advances in Genetic Programming, Vol. 1. MIT Press, Cambridge, pp. 355–368.

Arifovic, J., 1994. Genetic algorithm learning and the cobweb model. Journal of Economic Dynamics and Control 18 (1), 3–28.

Arifovic, J., 1995. Genetic algorithms and inflationary economies. Journal of Monetary Economics 36, 219–243.

Arifovic, J., 1996. The behavior of the exchange rate in the genetic algorithm and experimental economies. Journal of Political Economy 104 (3), 510–541.

Arifovic, J., Bullard, J., Duffy, J., 1997. The transition from stagnation to growth: an adaptive learning approach. Journal of Economic Growth 2, 185–209.

Arthur, W.B., Holland, J., LeBaron, B., Palmer, R., Tayler, P., 1997. Asset pricing under endogenous expectations in an artificial stock market. In: Arthur, W.B., Durlauf, S., Lane, D. (Eds.), The Economy as an Evolving Complex System II. Addison-Wesley, Reading, MA, pp. 15–44.

Bullard, J., Duffy, J., 1998a. A model of learning and emulation with artificial adaptive agents. Journal of Economic Dynamics and Control 22, 179–207.

Bullard, J., Duffy, J., 1998b. Learning and the stability of cycles. Macroeconomic Dynamics 2, 22–48.

Bullard, J., Duffy, J., 1999. Using genetic algorithms to model the evolution of heterogeneous beliefs. Computational economics 13, 41–60.

Chen, S.-H., Yeh, C.-H., 1996. Genetic Programming Learning and the Cobweb Model, Advances in Genetic Programming, Vol. 2. MIT Press, Cambridge, MA, 1996, pp. 443–466. (Chapter 22).

Chen, S.-H., Yeh, C.-H., 1997. Modeling speculators with genetic programming. In: Angeline, P., Reynolds, R.G., McDonnell, J.R., Eberhart, R. (Eds.), Evolutionary Programming VI. Lecture Notes in Computer Science, Vol. 1213. Springer, Berlin, pp. 137–147.

Chen, S.-H., Yeh, C.-H., 1998. Genetic programming in the overlapping generations model: an illustration with dynamics of the inflation rate. In: Porto, V.W., Saravanan, N., Waagen, D., Eiben, A.E. (Eds.), Evolutionary Programming VII. Lecture Notes in Computer Science, Vol. 1447. Springer, Berlin, pp. 829–838.

Chen, S.-H., Kuo, T.-W., 1999. Towards an agent-based foundation of financial econometrics: an approach based on genetic-programming artificial markets, In: Banzhaf, W., Daida, J., Eiben, A.E., Garzon, M.H., Honavar, V., Jakiela, M., Smith, R.E. (Eds.), GECCO-99: Proceedings of the Genetic and Evolutionary Computation Conference, July 13–17, 1999, Orlando, FL, USA, Vol. 2. Morgan Kaufmann, San Francisco, CA, pp. 966–973.

Chen, S.-H., Yeh, C.-H., 2000. Issues in the sensitivity analysis of agent-based modeling of artificial stock markets. AI-ECON Research Group Working Paper 2000-1, Department of Economics, National Chengchi University, Taipei, Taiwan.

Chen, S.-H., Duffy, J., Yeh, C.-H., 1996. Genetic programming in the coordination game with a chaotic best-response function. In: Angeline, P., Back, T., Fogel, D. (Eds.), Evolutionary Programming V. MIT Press, Cambridge, MA, pp. 277–286.

Chen, S.-H., Liao, C.-C., Yeh, C.-H., 2000a. On AIE-ASM: a software to simulate artificial stock markets with genetic programming. In: Chen, S.-H. (Ed.), Evolutionary Computation in Economics and Finance. Physica-Verlag, Wurzburg, forthcoming.

Chen, S.-H., Lux, T., Marchesi, M., 2000b. Testing for non-linear structure in an artificial financial market. Journal of Economic Behavior and Organization, forthcoming.

Dickey, D.A., Fuller, W.A., 1981. Likelihood ratio statistics for autoregressive time series with a unit root. Econometrica 49, 1057–1072.

Grossman, S.J., Stiglitz, J., 1980. On the impossibility of informationally efficiency markets. American Economic Review 70, 393–408.

Harrald, P., 1998. Economics and evolution. The panel paper given at the Seventh International Conference on Evolutionary Programming, March 25–27, San Diego, U.S.A.

Herreiner, D., 1998. Who learns from what from whom? Paper presented at the Fourth International Conference of the Society for Computational Economics, June 30–July 2, Cambridge, U.K.

Heymann, D., Pearzzo, R.P.J., Schuschny, A., 1998. Learning and contagion effects in transitions between regimes: some schematic multi-agents models. Journal of Management and Economics 2 (6).

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

Koza, J., 1992. Genetic Programming: On the Programming of Computers by means of natural selection, MIT.

LeBaron, B., 2000. Agent-based computational finance: suggested readings and early research. Journal of Economic Dynamics and Control 24, 679–702.

Lux, T., 1995. Herd behavior, bubbles and crashes. Economic Journal 105, 881–896.

Lux, T., 1998. The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distribution. Journal of Economic Behavior and Organization 33, 143–165.

Lux, T., Marchesi, M., 1999. Scaling and criticality in a stochastic multi-agent model of a financial market, Nature, 397, S. 498–500.

Miller, J.H., 1996. The coevolution of automata in the repeated prisoner's dilemma. Journal of Economic Behavior and Organization 29, 87–112.

Minsky, M.L., 1986. The Society of Mind. Simon & Schuster, New York.

Pagan, A., 1996. The econometrics of financial markets. Journal of Empirical Finance 3, 15–102.

Palmer, R.G., Arthur, W.B., Holland, J.H., LeBaron, B., Tayler, P., 1994. Artificial economic life: a simple model of a stock market. Physica D 75, 264–274.

Peters, E.E., 1991. A chaotic attractor for the S&P 500. Financial Analysts Journal, March/April.

Price, T.C., 1997. Using co-evolutionary programming to simulate strategic behavior in markets. Journal of Evolutionary Economics 7, 219–254.

Sargent, T.J., 1993. Bounded Rationality in Macroeconomics. Oxford.

Staudinger, S., 1998. Money as medium of exchange: an analysis with genetic algorithms. In: Proceedings of Conference on Computation in Economics, Finance, and Engineering, June 29–July 1, University of Cambridge.

Tayler, P., 1995. Modeling artificial stock markets using genetic algorithms. In: Goonatilake, S., Treleaven, P. (Eds.), Intelligent Systems for Finance and Business, pp. 271–288.

Vila, X., 1997. Adaptive artificial agents play a finitely repeated discrete principal-agent game. In: Conte, R., Hegselmann, R., Terna, P. (Eds.), Simulating Social Phenomena. Lecture Notes in Economics and Mathematical Systems, Springer, Berlin, pp. 437–456.

Vriend, N., 1998. An illustration of the essential difference between individual and social learning, and its consequence for computational analysis. Technical Report, Queen Mary and Westfield College, University of London, U.K.