

Genetic Programming in the Agent-Based Artificial Stock Market

Shu-Heng Chen
AI-ECON Research Group
Department of Economics
National Chengchi University
Taipei, Taiwan 11623
TEL: 886-2-27386874
FAX: 886-2-29361913
E-mail: chchen@nccu.edu.tw

Chia-Hsuan Yeh
AI-ECON Research Group
Department of Finance
I-Shou University
Kaohsiung County, Taiwan 84008
TEL: 886-7-6563711 ext 5718
FAX: 886-7-3552758
E-mail: spockyeh@ms38.hinet.net

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

Key Words: Agent-Based Computational Economics, Social Learning, Genetic Programming, Business School, Artificial Stock Markets, Simulated Annealing, Peer Pressure.

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). 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. MGA/MGP, in contrast, represents each agent as a *society of minds*. 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. 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. 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”.

Although Harrald’s criticism is well-acknowledged, we have seen no solution proposed to tackle this issue yet. 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. Therefore, if we supplement SGA/GP 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.

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. 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”. 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.

The rest of the paper is organized as follows. In Section 2, we shall present the analytical model on which our artificial market is constructed. A concrete application of genetic programming to the artificial stock market is also detailed. Experiment results and econometric analyses of these results are given in Section 3 followed by concluding remarks in Section 4.

2 The Analytical Model

The basic framework of the artificial stock market considered in this paper is the standard asset pricing model (Arthur et al., 1997). 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

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,

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

where $W_{i,t}$ is the wealth of trader i at time period t , and λ is the degree of relative risk aversion. 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.,

$$W_{i,t} = M_{i,t} + P_t h_{i,t} \quad (2)$$

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

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

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,

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

subject to

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

where $E_{i,t}(\cdot)$ 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*:

$$h_{i,t}^* = \frac{E_{i,t}(P_{t+1} + D_{t+1}) - (1+r)P_t}{\lambda\sigma_{i,t}^2}, \quad (6)$$

where $\sigma_{i,t}^2$ 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 $E_{i,t}(\cdot)$, 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

$$b_{i,t} = \begin{cases} h_{i,t}^* - h_{i,t-1}, & h_{i,t}^* \geq h_{i,t-1}, \\ 0, & \text{otherwise.} \end{cases} \quad (7)$$

and

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

Furthermore, let

$$B_t = \sum_{i=1}^N b_{i,t}, \quad O_t = \sum_{i=1}^N o_{i,t} \quad (9)$$

be the totals of the bids and offers for the stock at time t . Following Palmer et al (1994), we use the following simple rationing scheme:

$$h_{i,t} = \begin{cases} h_{i,t-1} + b_{i,t} - o_{i,t}, & \text{if } B_t = O_t, \\ h_{i,t-1} + \frac{O_t}{B_t}b_{i,t} - o_{i,t}, & \text{if } B_t > O_t, \\ h_{i,t-1} + b_{i,t} - \frac{B_t}{O_t}o_{i,t}, & \text{if } B_t < O_t. \end{cases} \quad (10)$$

All these cases can be subsumed into

$$h_{i,t} = h_{i,t-1} + \frac{V_t}{B_t}b_{i,t} - \frac{V_t}{O_t}o_{i,t} \quad (11)$$

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* $B_t - O_t$:

$$P_{t+1} = P_t(1 + \beta(B_t - O_t)) \quad (12)$$

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:

$$\beta(B_t - O_t) = \begin{cases} \tanh(\beta_1(B_t - O_t)) & \text{if } B_t \geq O_t, \\ \tanh(\beta_2(B_t - O_t)) & \text{if } B_t < O_t \end{cases} \quad (13)$$

where \tanh is the *hyperbolic tangent function*:

$$\tanh(x) \equiv \frac{e^x - e^{-x}}{e^x + e^{-x}} \quad (14)$$

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.,

$$H_t = \sum_i h_{i,t} = H. \quad (15)$$

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

$$M_{t+1} = \sum_i M_{i,t+1} = M_t(1+r) + H_t D_{t+1}. \quad (16)$$

2.3 Model of Adaptive Traders

Motivated by the *martingale hypothesis* in finance, we shall assume the following function form for $E_{i,t}(\cdot)$.

$$E_{i,t}(P_{t+1} + D_{t+1}) = (P_t + D_t)(1 + \theta_1 \tanh(\theta_2 \cdot f_{i,t})) \quad (17)$$

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 $\{i \mid f_{i,t} = 0\}$, denoted by $N_{1,t}$, we can know *how well the efficient market hypothesis is accepted among traders*. 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, originally we followed Arthur et al. (1997) to use the following updating scheme .

$$\sigma_{i,t}^2 = (1 - \theta_3)\sigma_{i,t-1}^2 + \theta_3[(P_t + D_t - E_{i,t-1}(P_t + D_t))^2]. \quad (18)$$

Without further restrictions, this update makes the subjective measure risk range between 0 and infinity. Since

$$h_{i,t}^* = \frac{E_{i,t}(P_{t+1} + D_{t+1}) - (1+r)P_t}{\lambda\sigma_{i,t}^2}, \quad (19)$$

it is clear that

$$h_{i,t}^* = \begin{cases} 0, & \text{if } \sigma_{i,t}^2 = \infty, \\ \pm\infty, & \text{if } \sigma_{i,t}^2 = 0. \end{cases} \quad (20)$$

As a consequence, all traders will tend to "leave" the market ($h_{i,t}^* = 0, \forall i$) due to incredible large subjective risks. However, the phenomenon characterized by

$$h_{i,t} \rightarrow 0, \text{ and } \sigma_{i,t}^2 \rightarrow \infty \quad (21)$$

is *not* self-fulfilling, because the increasing sequence of $\sigma_{i,t}^2$ implies a decreasing sequence of P_t .

If P_t is continuously decreasing, there is not much volatility and uncertainty, and traders have no reason to be subject to increasing subjective risks. Therefore, Equation (18) is not directly applicable to our model, and we have modified it into the following form,

$$\sigma_{i,t}^2 = (1-\theta_3)\sigma_{t-1|n_1}^2 + \theta_3[(P_t + D_t - E_{i,t-1}(P_t + D_t))^2]. \quad (22)$$

where

$$\sigma_{t|n_1}^2 = \frac{\sum_{j=0}^{n_1-1} [P_{t-j} - \bar{P}_{t|n_1}]^2}{n_1 - 1} \quad (23)$$

and

$$\bar{P}_t^{n_1} = \frac{\sum_{j=0}^{n_1-1} P_{t-j}}{n_1} \quad (24)$$

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

2.4 Business School and Single-Population GP

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. The faculty will be evaluated with a pre-specified 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 $GP_{i,t-1}$, 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 apply Schwefel's *1+1 strategy* 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 outcome of this *1+1* competition.

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 $gp_{i,t}$.

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 $\Delta W_{i,t}^{n_2}$ be this net change of wealth of trader i at time period t , i.e.,

$$\Delta W_{i,t}^{n_2} \equiv W_{i,t} - W_{i,t-n_2}, \quad (25)$$

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

$$p_{i,t} = \frac{R_{i,t}}{N}. \quad (26)$$

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

$$p_{i,t} < p_{j,t}, \text{ if } R_{i,t} < R_{j,t}. \quad (27)$$

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

$$\delta_{i,t}^{n_2} = \frac{W_{i,t} - W_{i,t-n_2}}{W_{i,t-n_2}}, \quad (28)$$

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

$$q_{i,t} = \frac{1}{1 + \exp^{\delta_{i,t}^{n_2}}}. \quad (29)$$

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

$$\lim_{\delta_{i,t}^{n_2} \rightarrow \infty} q_{i,t} = 0, \quad \lim_{\delta_{i,t}^{n_2} \rightarrow -\infty} q_{i,t} = 1. \quad (30)$$

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.

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 $gp_{j,t}$ ($j = 1, 2, \dots, 500$.) The class-taking behavior of traders is assumed to follow the following sequential search process. The trader will randomly select one *forecasting model* $gp_{j,t}$ ($j = 1, \dots, F$) 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.

3 Simulation Results

Based on the experiment 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 while deriving the price P_t . In the following, we shall present our results in an order to answer a series of questions raised in Pagan (1996).¹

First, are price and returns normally distributed? The time series plot of the stock price is drawn in Figure 1. Over this long horizon, P_t fluctuates between 55 and 105. The basic statistics of this series, $\{P_t\}_{t=1}^{14000}$, is summarized in Table 2. Given the price series, the return series is derived as usual,

$$r_t = \ln(P_t) - \ln(P_{t-1}). \quad (31)$$

Figure 2 is a time series of stock return, and Table 2 also gives the basic statistics of this return series. From Table 2, 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 especially striking in the return series.

Second, does prices follow a random walk? Or, more technically, 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.

$$\Delta \ln(P_t) = \ln(P_t) - \ln(P_{t-1}) = \beta_1 \ln(P_{t-1}) \quad (32)$$

¹Pagna (1996) summarized a list of *stylized facts* in financial time series.

Table 1: Parameters of the Stock Market

The Stock Market	
Shares of the stock (H)	100
Initial money supply (M_1) per capital	100
Interest rate (r)	0.1
Stochastic Process (D_t)	$U(5.01, 14.99)$
Price adjustment function	<i>tanh</i>
Price adjustment (β_1)	10^{-5}
Price adjustment (β_2)	0.2×10^{-5}
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	$\{+, -, Sin, Cos, Exp, Rlog, Abs, Sqrt\}$
Terminal set	$\{P_t, P_{t-1}, \dots, P_{t-10}, P_{t-1} + D_{t-1}, \dots, P_{t-10} + D_{t-10}\}$
Selection scheme	Tournament selection
Tournament size	2
Probability of creating a tree by reproduction	0.10
Probability of creating a tree by crossover	0.70
Probability of creating a tree by mutation	0.20
Probability of mutation	0.0033
Probability of leaf selection under crossover	0.5
Mutation scheme	Tree Mutation
Replacement scheme	(1+1) Strategy
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 $\sigma_{t n_1}^2$ (n_1)	10
Evaluation cycle (n_2)	1
Sample size (n_3)	10
Search intensity (I^*)	5
θ_1	0.5
θ_2	10^{-5}
θ_3	0.0133

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, 3, 4, 5, and 6. For details, see Koza (1992).

Table 2: Basic Statistics of the Artificial Stock Price Series and Return Series

Price Series						
Periods	\bar{P}	σ	SK	KU	JB	p
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
Return Series						
\times	10^{-4}	10^{-2}	1	1	10^3	1
1-2000	-0.74	1.5	3.53	23.64	39.67	0.00
2001-4000	-0.57	1.0	3.26	18.83	24.46	0.00
4001-6000	-0.18	0.7	3.72	25.94	48.48	0.00
6001-8000	-0.24	0.7	3.70	25.79	47.86	0.00
8001-10000	0.32	0.7	3.69	26.97	52.45	0.00
10001-12000	1.69	1.0	6.91	86.56	597.87	0.00
12001-14000	-1.54	0.9	4.18	32.80	79.86	0.00

“SK” refers to skewness, “KU” refers to kurtosis, “JB” Jarqu-Bera test, and “p” p-value.

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 3, from the total number of 7 periods none leads to a rejection of the presence of a unit root. All of this does suggest that P_t does follow a random walk.

Third, are returns independently and identically distributed? Here, we followed the procedure of Chen and Kuo (1999). 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 3

Table 3: Unit Root Test and PSC Filtering

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)

The MacKinnon critical values for rejection of hypothesis of a unit root at 99% (95%) significance level is -2.5668 (-1.9395).

Table 4: BDS Test

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

The test statistic is asymptotically normal with mean 0 and standard deviation 1. The significance level of the test is set at 0.95.

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, and the result is given in Table 4. 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 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

$$E_t(P_{t+1} + D_{t+1}) = P_t + D_t? \quad (33)$$

If they do not believe in the martingale hypothesis, do they search intensively? To answer the first question, the time series of the number of martingale believers, $N_{1,t}$, is drawn in Figure 3. The figure is drawn only up to the first 2000 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.

This naturally brings up the second question: *if they do not believe in the martingale hypothesis, what do they actually do?* Figure 4 is the time

Table 5: Microstructure Statistics: Average of Traders with Successful Search and Complexity of Evolving Strategies

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

\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.

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 5. From Table 5, 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 2, what business school offers is a collection of forecasting models $\{gp_{i,t}\}$, which can 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 this long-run property. What motivates them to search and helps them to survive is in effect *brief signals*.

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.

$$k_t = \frac{\sum_i^{500} k_{i,t}}{500}, \quad \text{and} \quad \kappa_t = \frac{\sum_i^{500} \kappa_{i,t}}{500}. \quad (34)$$

Figures 5 and 6 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 (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 (Table 5). Despite the rejection of the monotone hypothesis, we see no evidence that traders' behavior will converge to the simple martingale model.

4 Concluding Remarks

In this paper, we propose a new architecture of the artificial stock market. The single-run simulation with 14,000 trading days can, at best, be considered as a pilot experiment. But, from this pilot experiment, we already experienced the rich dynamics generated from the agent-based modeling of the artificial stock market.

Acknowledgement

Research support from NSC grants No.88-2415-H-004-008 and No. 87-I-2911-15 is gratefully acknowledged. The authors are also grateful for the helpful comments from three anonymous referees.

References

- [1] Arthur, W. B., J. Holland, B. LeBaron, R. Palmer and P. Tayler (1997), "Asset Pricing under Endogenous Expectations in an Artificial Stock Market," in W. B. Arthur, S. Durlauf & D. Lane (eds.), *The Economy as an Evolving Complex System II*, Addison-Wesley, pp. 15-44.
- [2] Chen, S.-H. and T.-W. Kuo (1999), "Towards an Agent-Based Foundations of Financial Econometrics: An Approach Based on Genetic-Programming Artificial Markets," in W. Banzhaf et al. (eds), *GECCO-99: Proceedings of the Genetic and Evolutionary Computation Conference*, San Francisco, CA: Morgan Kaufmann.
- [3] Dickey, D. A. and W. A. Fuller (1981), "Likelihood Ratio Statistics for Autoregressive Time Series with a Unit Root," *Econometrica* 49, pp. 1057-1072.
- [4] Grossman, S. J. and J. Stiglitz (1980), "On the Impossibility of Informationally Efficiency Markets," *American Economic Review*, 70, pp. 393-408.
- [5] Harrauld, 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.
- [6] Pagan, A. (1996), "The Econometrics of Financial Markets," *Journal of Empirical Finance*, 3, pp. 15-102.

