Journal of Economic Behavior & Organization
Vol. 49 (2002) 217–239

JOURNAL OF
Economic Behavior
& Organization

www.elsevier.com/locate/econbase

On the emergent properties of artificial stock markets: the efficient market hypothesis and the rational expectations hypothesis[☆]

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

[a] *Department of Economics, AI-ECON Research Center, National Chengchi University, Taipei 11623, Taiwan*
[b] *Department of Information Management, Yuan Ze University, Chungli, Taoyuan 320, Taiwan*

Received 9 April 2001; received in revised form 3 August 2001; accepted 9 August 2001

Abstract

By studying two well known hypotheses in economics, this paper illustrates how emergent properties can be shown in an agent-based artificial stock market. The two hypotheses considered are the efficient market hypothesis and the rational expectations hypothesis. We inquire whether the macrobehavior depicted by these two hypotheses is consistent with our understanding of the microbehavior. In this agent-based model, genetic programming is applied to evolving a population of traders learning over time. We first apply a series of econometric tests to show that the EMH and the REH can be satisfied with some portions of the artificial time series. Then, by analyzing traders' behavior, we show that these aggregate results cannot be interpreted as a simple scaling-up of individual behavior. A conjecture based on sunspot-like signals is proposed to explain why macrobehavior can be very different from microbehavior. We assert that the huge search space attributable to genetic programming can induce sunspot-like signals, and we use simulated evolved complexity of forecasting rules and Granger causality tests to examine this assertion. © 2002 Elsevier Science B.V. All rights reserved.

JEL classification: G12: asset pricing; G14: information and market efficiency; D83: search, learning, and information

Keywords: Artificial stock markets; Emergent properties; Efficient market hypothesis; Rational expectations hypothesis; Genetic programming

☆ This paper is modified and extended from an earlier version presented at the 4th Workshop on Economics with Heterogeneous Interacting Agents (WEHIA'99), Genoa, Italy, 4–5 June 1999.

* Corresponding author.
E-mail addresses: chchen@nccu.edu.tw (S.-H. Chen), imcyeh@saturn.yzu.edu.tw (C.-H. Yeh).

PII: S0167-2681(02)00068-9

1. Background and motivation

While it is claimed quite frequently that the stock market is a complex adaptive system, conventional financial models, constrained by computing power, are not capable of demonstrating this feature. However, recent progress in computing technology has made possible a more ambitious vehicle to construct and simulate the stock market. The fledgling research field, known as the artificial stock market, is distinguished from the conventional model-building in many essential ways.[1] Generally speaking, models in this field are composed of many heterogeneous interacting adaptive traders. The conventional devices such as the rational expectations hypothesis and the representative agent are discarded (Arthur, 1992). In principle, the artificial stock market is a promising way to study the stock market as a complex adaptive system. By that, we mean two things. First, the artificial stock market is rich in dynamics. Second, it is rich in emergent properties.[2]

The rich dynamics of the artificial stock market have been documented in the literature. One of the early attempts of this research was to show that many econometric properties (stylized facts) of financial time series can be replicated by artificial stock markets. The properties replicated include volatility clustering (autoregressive conditional heteroskedasticity (ARCH)), excess kurtosis (fat-tail distribution), bubbles and crashes, chaos, unit roots, and many others.[3] Thus, there is little doubt that the artificial stock market can generate rich dynamics. However, being able to generate rich dynamics is only a minor part of complex adaptive systems. To be a complex adaptive system, rich dynamics must be generated endogenously (or from bottom up), rather than be given exogenously (top down). It is this difference that leads to the main characteristic of complex adaptive systems, namely, emergence.

Emergence is about "how large interacting ensembles exhibit collective behavior that is very different from anything one may have expected from simply scaling up the behavior of the individual units" (Krugman, 1996, p. 3), or ". . . in a structured system, new properties emerge at higher levels of integration which could not have been predicted from a knowledge of the lower level components" (Mayr, 1997, p. 19). Examples of emergence abound in other fields (Holland, 1998), and economists are anything but unfamiliar with the significance of this term. Apart from the Santa Fe Institute Economists, Krugman (1996) and Epstein and Axtell (1996) are among the first few economists who exemplified emergence with a series of economic phenomena. Nevertheless, the emergent properties of the artificial stock market has not received a full attention. This paper considers a different research direction. Instead of replicating the econometric properties of financial time series, though it is still worth doing, we are concerned with identifying some areas of the artificial stock market where the phenomena observed can be plausibly argued as emergent behavior. The areas considered in this paper are two celebrated hypotheses in economics and finance, namely, the efficient market hypothesis (EMH) and the rational expectations hypothesis (REH).

[1] See LeBaron (2000).

[2] Publishing the article "More is Different" in 1972s Science magazine, Philip Anderson, the 1977 Nobel laureate physicist, may be regarded as the father of the science of emergence.

[3] See, e.g. Lux (1995, 1997, 1998), Lux and Marchesi (1999), and LeBaron et al. (1999).

First, the efficient market hypothesis. What does it mean if the EMH can be an emergent property? Consider the EMH as a collective behavior. It would be an emergent property if it is not expected from our understanding of the behavior of individual traders. Let us take an extreme case. Suppose that none of the traders believe in the EMH, then this property will not be expected to be a feature of their collective behavior. Even if it is observed, it has no direct link to the individual behavior. So, if the collective behavior of these traders indeed satisfies the EMH as tested by the standard econometric procedures, then we would consider the EMH as an emergent property. Second, the rational expectations hypothesis. Consider the rational expectations hypothesis as a collective behavior. It would be an emergent property if all our traders are boundedly rational, with their collective behavior satisfying the REH as tested by econometrics.

This way of identifying the EMH and the REH as emergent properties may not seem rigorous, since the word "expected" or "surprising" is somewhat subjective. Nonetheless, in the light of the long-lasting debate on the two hypotheses, this particular way of defining emergent properties is quite normal, if not satisfactory. In a sense, it provides a new perspective to reflect upon these controversies. Consider a spectrum. On the leftmost are individuals, and the rightmost an aggregate of individuals (the representative agent). It seems much easier to reject these hypotheses at the leftmost point than at the other extreme. To the leftmost is the area of psychology, where evidence of bounded rationality dates back to 1970s (Tversky and Kahneman, 1974), whereas to the rightmost is the turf of the representative agent whose rational behavior has been evidenced by advanced econometrics since Hall's work (1978) on consumption theory in 1978.

Therefore, instead of thinking of this spectrum as an encapsulation of conflicting viewpoints, one may consider it a system whose microbehavior is rather different from macrobehavior (Kirman, 1992). Of course, this manner of thinking is nothing new in economics. A list of early examples, such as Adam Smith's invisible hand and Hayek's hypothesis, was well documented in Krugman (1996). However, what has not been done in conventional economics is to construct a system (artificial society) where to allow both views of the world are represented. What we have in mainstream economics is a highly abstract representative agent. Under these circumstances, there is no distinction between the microbehavior and the macrobehavior of traders, and hence, no room for the study of emergent properties. Recent advancement in research technology provides us with an opportunity to address the issue. This paper is not the first one in this line of research, and certainly not the last. What distinguishes this study from earlier ones, however, is that this study may be viewed as a first attempt to formally interpret the EMH and the REH as emergent properties in the context of artificial stock markets.

The rest of the paper is organized as follows. In Section 2, we present the analytical model upon which our artificial stock market is built. Genetic programming is introduced in Section 3 to model a population of traders learning over time. We then present and analyze the results of simulations in Sections 4 and 5. A series of econometric tests and microstructure analysis were conducted to show that the EMH and REH can be emergent properties which will be further discussed in Section 6. Concluding remarks are given in Section 7.

2. The artificial stock market

2.1. 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 traders each of them has the goal to maximize her expected utility based on her forecast of the future. Technically, there are two major components of this market, namely, traders and their interactions.

2.2. 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 in. 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 $h_{i,t}$ denote 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 stationary stochastic process. In this paper, we assume that D_t is an i.i.d. normal process with mean μ and variance σ_ξ^2; r is the riskless interest rate. Given this wealth dynamic, 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 Eq. (3), where $E_{i,t}(\cdot)$ is trader i's conditional expectations of W_{t+1} given her information up to t (the information set $I_{i,t}$). The choice variable of this optimization problem is $h_{i,t}$.

It is well known that under CARA utility and Gaussian distribution for forecasts, trader i's desire demand, $h^*_{i,t}$, for holding shares of the risky asset is linear in the expected excess return (Grossman and Stiglitz, 1980, p. 396):

$$h^*_{i,t} = \frac{E_{i,t}(P_{t+1} + D_{t+1}) - (1+r)P_t}{\lambda \sigma^2_{i,t}}, \tag{5}$$

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 $E_{i,t}(\cdot)$, which will be given in detail later.

2.3. 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 buy at period t, and let $o_{i,t}$ be the number of shares trader i would like to sell at period t. It is clear that

$$b_{i,t} = \begin{cases} h^*_{i,t} - h_{i,t-1}, & h^*_{i,t} \geq h_{i,t-1}, \\ 0, & \text{otherwise} \end{cases} \tag{6}$$

and

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

Clearly, from Eqs. (6) and (7), the trader can only buy or sell but not both at the same time.

Furthermore, let

$$B_t = \sum_{i=1}^{N} b_{i,t} \tag{8}$$

and

$$O_t = \sum_{i=1}^{N} o_{i,t} \tag{9}$$

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:

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

All these cases can be subsumed into

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

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 $B_t - O_t$:

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

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

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

where tanh is the hyperbolic tangent function.

2.4. *Model of adaptive traders*

In this section, we will address the formation of traders' expectations, $E_{i,t}(P_{t+1} + D_{t+1})$ and $\sigma^2_{i,t}$. There are several ways to generate $E_{i,t}(\cdot)$. First, $E_{i,t}(\cdot)$ can be a function directly generated by genetic programming. This is certainly the most straightforward way. But, in this case, the range of $E_{i,t}(\cdot)$ can hardly be held in check, and that can potentially make the price dance in a crazy manner. Second, instead of level, one may be interested in forecasting increments instead of level. That is what we are doing in this paper. The specific function form we consider is the following,

$$E_{i,t}(P_{t+1} + D_{t+1}) = (P_t + \mu)(1 + \theta_1 \tanh(\theta_2 f_{i,t})). \tag{14}$$

where $f_{i,t}$ generated by genetic programming is fed into a hyper-tangent transformation, and that restricts the range to a $(-1, 1)$ open interval.[4] After being pre-multiplied by a constant θ_1, its range is $(-\theta_1, \theta_1)$. So, Eq. (14) can be interpreted as a forecast of the growth rate (rate of return) of $P_t + \mu$.

Apart from the technical reason given above, function (14) has two other advantages in economic sense. First, in the case of homogeneous rational expectations equilibrium, the price is a constant P^*, and the best forecast (conditional expectations) is just[5]

$$E_t(P_{t+1} + D_{t+1}) = P_t + \mu = P^* + \mu. \tag{15}$$

In terms of Eqs. (14), (15) is a special case corresponding to $f_{i,t} = 0$. Furthermore, since D_t follows an i.i.d. process with mean μ, it would be interesting to consider the case where P_t fails to converge to P^*, but instead follows a random walk. In this case,

$$E_t(P_{t+1} + D_{t+1}) = P_t + \mu, \tag{16}$$

or simply $E_t(P_{t+1}) = P_t$.[6] In plain English, the best forecast for tomorrow's price is today's price, a property known as the martingale hypothesis in finance. Again, Eq. (16) is a special case of Eq. (14) when $f_{i,t} = 0$.

So, when the condition for Eq. (16) holds, i.e. when the martingale hypothesis is validated as an aggregate phenomenon, $f_{i,t} = 0$ would imply that macrobehavior is consistent with microbeliefs. Alternatively speaking, on the top the market is efficient in the martingale

---

[4] The population of functions $f_{i,t}$ ($i = 1, \ldots, N$) is determined by the genetic programming procedure to be detailed in the following section.

[5] See Arthur et al. (1997).

[6] More rigorously, one needs the assumption that D_t fails to Granger cause P_t. But, theoretically, this is self-evident because the sequence $\{D_t\}$ is independent. For technical details, see Arthur et al. (1997).

sense, and on the bottom trader i also believes so. Therefore, from the cardinality of the set $\{i \mid f_{i,t} = 0\}$, denoted by $N_{1,t}$, one can know how well the efficient market hypothesis is accepted among traders, or how well macrobehavior is consistent with microbeliefs.

As to the subjective risk equation, we adopt a modified version of the equation originally used by Arthur et al. (1997):

$$\sigma^2_{i,t} = (1-\theta_3)\,\sigma^2_{t-1|n_1} + \theta_3[(P_t + D_t - E_{i,t-1}(P_t + D_t))^2], \tag{17}$$

where

$$\sigma^2_{t|n_1} = \frac{\sum_{j=0}^{n_1-1}[P_{t-j} - \bar{P}_{t|n_1}]^2}{n_1 - 1}, \quad \text{and} \quad \bar{P}_{t|n_1} = \frac{\sum_{j=0}^{n_1-1} P_{t-j}}{n_1}. \tag{18}$$

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

3. Modeling traders' adaptation with genetic programming

In this paper, we will use genetic programming to evolve traders' forecasts denoted by $f_{i,t}$. Like genetic algorithms, genetic programming is a model for a population of agents learning over time. In this paper, it is used to evolve a population of forecasting function $\{f_{i,t}\}$ held by traders, as outlined in the following sections.

First, each trader is working under survival pressure. The pressure traders have to bear prompts them to search (struggle) for something better, in our case a better forecasting rule. The stronger the pressure, the stronger the incentive to search. Second, once a trader decides to search, the whole search process is driven by genetic programming. On the other hand, if a trader decides not to search, her forecasting rule shall remain unchanged.

3.1. *Pressure*

Pressure is basically a psychological term which is sometimes not very objective and hence not very easy to model. What we propose here is to quantify two basic types of pressure well studied by psychologists, namely, peer pressure and self-pressure. The way to quantify peer pressure is as follows. Suppose that traders will examine how well they have performed over the last n_2 trading days, when compared with other traders. Moreover, they rank their performance by the net change of wealth over the last n_2 trading days. Let $W^{n_2}_{i,t}$ be this net change of wealth over a n_2 time horizon of trader i at time period t, i.e.

$$\Delta W^{n_2}_{i,t} \equiv W_{i,t} - W_{i,t-n_2}, \tag{19}$$

and, let $R^{n_2}_{i,t}$ be her ranking. Then $R^{n_2}_{i,t}$ can be considered as a measure for peer pressure trader i bears at time t.

In addition to peer pressure, traders may also examine how much progress they have made over the last n_2 trading days, i.e. the growth rate of income over the last n_2 days:

$$\delta^{n_2}_{i,t} = \frac{\Delta W^{n_2}_{i,t} - \Delta W^{n_2}_{i,t-n_2}}{|\Delta W^{n_2}_{i,t}|}. \tag{20}$$

This measure $\delta^{n_2}_{i,t}$ defines the self-pressure for trader i.

3.2. Search incentive

Under peer pressure and self-pressure, trader i's incentive to search can be modeled as follows. First of all, trader i has a particular probability of searching. Denote this probability by $p_{i,t}$, and assume that

$$p_{i,t} = \frac{R_{i,t}^{n_2}}{N}. \tag{21}$$

The choice of function (21) is quite intuitive. It simply means that

$$p_{i,t} < p_{j,t}, \quad \text{if} \quad R_{i,t}^{n_2} < R_{j,t}^{n_2}. \tag{22}$$

In other words, traders who come out top shall suffer less peer pressure, and hence be less motivated to search than those who are ranked at the bottom.

In addition to $p_{i,t}$, under pressure $\delta_{i,t}^{n_2}$, there is another chance with which trader i will search. This probability, denoted by $q_{i,t}$, is assumed to be:

$$q_{i,t} = \frac{1}{1 + \exp(\delta_{i,t}^{n_2})}. \tag{23}$$

The choice of density function (23) is also straightforward. Notice that

$$\lim_{\delta_{i,t}^{n_2} \to \infty} q_{i,t} = 0, \quad \text{and} \quad \lim_{\delta_{i,t}^{n_2} \to -\infty} q_{i,t} = 1. \tag{24}$$

Therefore, the traders who have made great progress will naturally be more confident, and thus, have little incentive to search, whereas those who suffer devastating regression will have a strong desire to search.

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

$$r_{i,t} = p_{i,t} + (1 - p_{i,t})q_{i,t} = \frac{R_{i,t}^{n_2}}{N} + \frac{N - R_{i,t}^{n_2}}{N} \frac{1}{1 + \exp(\delta_{i,t}^{n_2})} \tag{25}$$

And the number of traders who decide to search, denoted by $N_{2,t}$, is a sum of these Bernoulli random variables:

$$N_{2,t} = \sum_{i=1}^{N} \chi_{i,t}, \tag{26}$$

where $\chi_{i,t}$ is a Bernoulli random variable with a probability of success $r_{i,t}$.

3.3. Search process

The search process to be detailed below will determine the outcome of the trader's search. Basically, it describes how promising ideas (forecasting rules) are popularized or how new ideas are discovered during traders' search and adaptive process, i.e. the evolution of a population of ideas. Recently, genetic algorithms (GAs) and genetic programming (GP) have been extensively used to substantiate processes like this. However, the straightforward application of single-population GAs or GP, which rests on the assumption that strategies are observable and imitable, has been seriously criticized by Harrald (2000). Harrald (2000) 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.

Chen and Yeh (2001) had a lengthy discussion on this issue, and proposed a solution to it. The solution is based on an idea called business school. They argued that, in the case of school, observability and imitability (replicatability) is not an assumption but a rule. Hence, there is no distinction between genotype and phenotype in school. Thus, the original operation of single-population GP in a society of traders can now be replaced by a society of faculty members. Single-population GP is then conducted in a standard way. In a separate study (Chen and Yeh, 2000), they also found that the time series generated by the architecture without the business school are highly serially correlated and linearly predictable, which casts doubts on whether running single-population GP directly on a society of traders is a proper design for agent-based financial markets. Accordingly, in this paper, we only present the architecture with business school.

In Chen and Yeh (2001), the business school mainly consists of faculty members. Let F be the number of faculty members (forecasting rules). Each faculty member (forecasting model) is represented by a tree (GP parse tree). The faculty will be evaluated with 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 rule (faculty member) will be reviewed by a visitor. The visitor is another model which is generated randomly from the collection of the existing rules 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 $p_m (= 1 - p_r - p_c)$. In the case of reproduction, we first randomly select two GP trees, say, $gp_{j,t-1}$ and $gp_{k,t-1}$. The mean absolute percentage error (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 the lower MAPE, say $gp_{j,t-1}$, is selected. We then run a tournament over the host $gp_{i,t-1}$ and the visitor $gp_{j,t-1}$ based on the criterion MAPE, and $gp_{i,t}$ is the winner of this tournament. In the case of mutation, we follow the same procedure as reproduction except that, before meeting its match $gp_{i,t-1}$, $gp_{j,t-1}$ has a chance of being perturbed to $gp'_{j,t-1}$ by tree mutation.

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 $gp_{i,t}$.

Given the business school described above, traders' decision to search is a decision to go to school. 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 F faculty members (forecasting rules), let us denote them by $gp_{j,t}$ $(j = 1, 2, \ldots, F)$. The class-taking behavior of traders is assumed to follow the following sequential search process. The trader will randomly select one forecasting rule $gp_{j,t}$ $(j = 1, \ldots, F)$ with a uniform distribution. She will then experiment 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. If this happens, the trader is considered to have a successful search. 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 number of successful searchers in time t is denoted by $N_{3,t}$.

4. The efficient market hypothesis

4.1. Macrobehavior

Based on the parameter values specified in Table 1, a single run with 20,000 trading days was conducted. This generated a time series of the artificial stock price with 20,000 observations. A time series plot of this artificial price series is given in Fig. 1. In addition, Table 2 gives the basic statistics of the return series, including the mean, the standard deviation, etc.[7] During the whole simulation period, the price ranges from 75 to 100 with a mean 85, which is very close to the homogeneous rational expectations equilibrium (HREE) price 80.[8] Fig. 2 exhibits the time series plot of the stock return $\{r_t\}$, where $r_t = \ln(P_t) - \ln(P_{t-1})$.[9] We can see that the return series centers around zero. However, the distribution is not symmetric. In particular, there seems to be a flat (Fig. 2) for negative returns, but not for positive returns. The flat is situated at -0.010050, i.e. the value of $\ln(1 + \beta(-500))$. This is the value corresponding to the case where all the traders would like to sell all their stocks, and none would like to buy. What happens then is an excess offer to sell up to 500 units. By the price adjustment equation (Eq. (13)), this sell pressure will lead to a return which is $\ln(1 + \beta(-500))$, i.e. -0.010050. Since traders are not allowed to sell short, the excess offer will not be higher than 500, and r_t cannot be lower than -0.010050. On the other hand, traders can buy whatever amount they can afford; hence, there is no such flat for positive returns.

The first macrobehavior we would like to examine is to see whether our artificial stock market is efficient in the sense that the stock returns are statistically independent. To do so,

[7] Most statistics reported in the following are based on a non-overlapping decomposition of the original 20,000 into 10 subperiods, each with 2000 observations.

[8] Under full information and homogeneous expectations, the HREE price is given in Eq. (27) (Arthur et al., 1997, pp. 40–41).

$$P_t = \frac{1}{r}\left[\mu - \lambda\sigma_\xi^2\left(\frac{H}{N}\right)\right] = \frac{1}{r}\left[\mu - \lambda\sigma_\xi^2 h\right]. \tag{27}$$

Since in our experiments (Table 1), $(\mu, \sigma_\xi^2, r, \lambda, h) = (10, 4, 0.1, 0.5, 1)$, the HREE price is 80.

[9] A detail analysis of the statistical behavior of the stock return can be found in Chen and Yeh (2000).

Table 1
Parameters of the stock market

The stock market	
Shares of the stock (H)	100
Initial money supply (M_1)	100
Interest rate (r)	0.1
Stochastic process (D_t)($\sim$ normal(μ, σ_ξ^2)	i.i.d. normal(10, 4)
Price adjustment function	tanh
Price adjustment (β_1)	0.2×10^{-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	$\{+, -, \times, /$, sin, cos, exp, rlog, abs, 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	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
Maximum depth of tree	17
Number of generations	20000
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

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).

we followed the procedure of Chen, Lux and Marchesi (2000). This procedure is composed of two steps, namely, the PSC filtering and the BDS testing. We first applied the Rissanen's (Rissanen, 1989) predictive stochastic complexity (PSC) to the return series. The PSC criterion is a model selection criterion. It selects the model with the minimum PSC. By the PSC criterion, we can identify the linear ARMA model (p, q) of a series. If a series satisfies

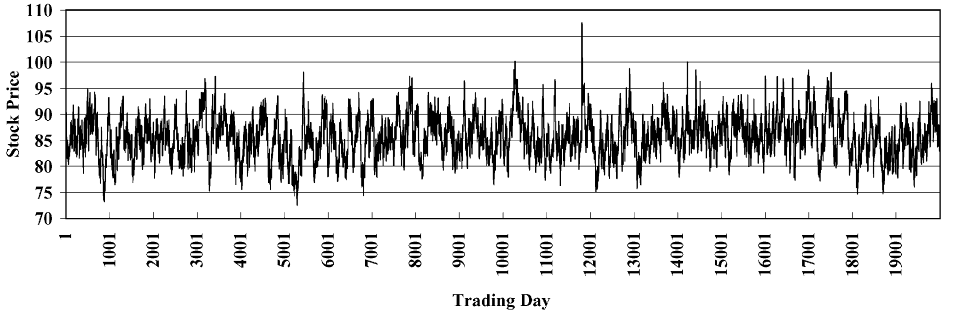

Fig. 1. Time series plot of the stock price.

Table 2
Basic statistics of the artificial stock return series and tests for i.i.d.

Periods	$\bar{r}$	σ	SKew	Kur	PSC	BDS	GARCH	Kaplan
1–2000	−4.17	9.54	1.16	5.01	(0, 0)	A (0.69)	R (0, 1)	A
2001–4000	−3.81	9.53	0.98	3.78	(0, 0)	A (1.74)	R (0, 1)	R
4001–6000	−0.27	9.63	1.24	5.69	(1, 0)	A (1.00)	R (0, 1)	R
6001–8000	3.92	9.78	1.20	4.94	(0, 0)	A (1.05)	R (0, 1)	A
8001–10000	1.24	9.81	1.37	6.68	(0, 0)	A (1.12)	A (0, 0)	R
10001–12000	0.01	10.94	5.15	92.67	(0, 0)	A (0.35)	A (0, 0)	R
12001–14000	−1.62	9.60	1.22	5.44	(0, 0)	A (1.10)	R (0, 1)	Am
14001–16000	2.78	10.61	2.76	26.28	(0, 0)	A (1.44)	A (0, 0)	R
16001–18000	−3.06	10.09	1.85	13.51	(0, 0)	A (1.15)	A (0, 0)	A
18001–20000	−0.01	9.41	0.87	3.49	(0, 0)	A (1.29)	R (0, 1)	R

Here, $\bar{r}$ and σ are 10^5 and 10^3 times of the respective original value. "SKew" refers to "skewness", and "Kur" refers to kurtosis. The pairs of numbers (p, q) in the column "PSC" are the orders of the ARMA(p, q) model selected by the PSC criterion. The test result "A" (accept) or "R" (reject) in the column "BDS" is based on a significance level at 0.05. Inside the bracket is the BDS test statistic $J_{\epsilon,m}$. The test result shown in the column "GARCH" are based on the Lagrange multiplier (LM) test. The pairs of numbers inside the bracket are the orders of the GARCH(p, q) model selected by the SIC criterion. "Am" shown in the Kaplan test refers to "ambiguous".

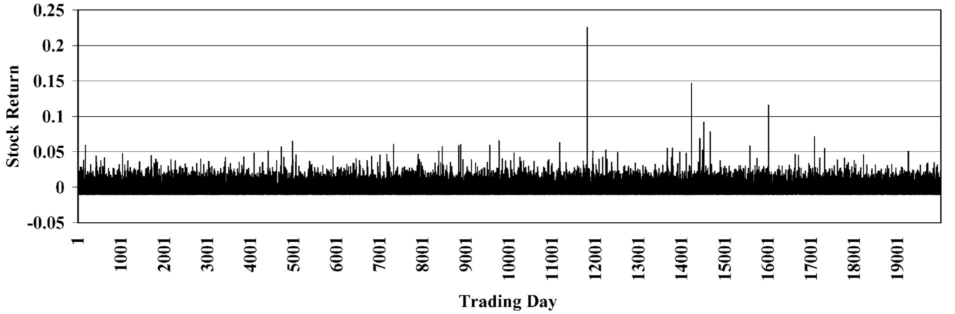

Fig. 2. Time series plot of the stock return.

the EMH, then both p and q should be 0, i.e. there is no linear dependence, and hence, the series is not linearly predictable. The fifth column of Table 2 gives us the ARMA(p, q) process extracted from the return series. Interestingly enough, all these seven periods are linearly uncorrelated ($p = 0, q = 0$) except the third period 4001–6000, which is identified as an AR(1) process.

We filtered out the linear signal detected so that any signal left in the residual series must be non-linear. One of the most frequently used tests for non-linear dependence is the celebrated BDS test (Brock et al., 1996). The BDS test is a test based on the fact that, if $\{r_t\}$ is an i.i.d. series, then the BDS test statistic has a limiting standard normal distribution. The BDS test was then applied to the prewhitened series. Since most return series have no linear signal, the BDS test was simply directly applied to the original series except the for third one. There are two parameters required to conduct the BDS test. One is the distance parameter (ϵ standard deviations), and the other the embedding dimension (DIM). Since our results are not sensitive to either choice, only those with $\epsilon = 1$ and DIM $= 5$ are reported here. The results are given in the sixth column of Table 2. By the BDS test, what is interesting is that the null hypothesis of i.i.d. (identically and independently distributed) is not rejected in any of these ten periods.

Motivated by Barnett et al. (1998), we also carried out the Lagrange multiplier (LM) test for the presence of the autoregressive conditional heteroskedasticity (ARCH) effect of the residual. We took lags up to 12. If the null hypothesis is rejected, we will further identify the GARCH (generalized ARCH) order of the series by the Schwartz Information Criterion (SIC). These results are exhibited in the seventh column of Table 2. Out of the 10 subperiods, six exhibit the ARCH effect. By the SIC criterion, they are all GARCH(0, 1). This result is somewhat inconsistent with the BDS test in non-linear independence. Here, six out of the 10 series which fail to reject the null of i.i.d. under the BDS test are now identified as GARCH(0, 1) series.

Finally, considering possibility that the time series may be chaotic rather than stochastic, we also conduct the Kaplan test (Kaplan, 1994). Kaplan (1994) used the fact that deterministic processes, unlike stochastic processes, have the following property: points that are nearby are also nearby under their image in phase space. That is, if X_i and Y_j are close to each other, then X_{i+1} and Y_{j+1} are also close to each other. Technically speaking, let $X_i = (r_i, r_{i-\tau}, r_{i-2\tau}, \dots, r_{i-(m-1)\tau})$ embedded in m-dimensional phase space, then there is a recursive function given

$$X_{i+\tau} = f(X_i) \tag{28}$$

with the fixed positive integer time delay τ. For a given choice of embedding dimension m, one can calculate

$$\delta_{ij} = |X_i - X_j|, \quad \text{and} \quad \epsilon_{i,j} = |X_{i+\tau} - X_{j+\tau}|, \tag{29}$$

for all pairs of time subscripts (i, j). Let $E(\zeta) = \sum_{A_\zeta} \epsilon_{i,j}/\#\{A_\zeta\}$, where $A_\zeta \equiv \{(i, j) : \delta_{i,j} < \zeta\}$. For a perfectly deterministic system with continuous f, one expects to have $\lim_{\zeta \to 0} E(\zeta) = 0$. Based on this theoretical property, Kaplan's value K is defined as the limit of $E(\zeta)$ as $\zeta \to 0$.

The essential ingredient of the Kaplan test is to find a piecewise regression line for $(\delta_{i,j}, \epsilon_{i,j})$ and use the intercept as an estimation of the Kaplan value K, $\hat{K}$. The test statistic for $\hat{K}$ is produced based on the linear surrogates, i.e. the artificial series which has the same histogram and a similar autocorrelation function as the target series. By this test statistic, one can then accept (reject) the null of i.i.d. if $\hat{K}$ is smaller (larger) relative to the test statistic. A Matlab computational algorithm prepared by Kaplan is used to generate $\hat{K}$. Barnett et al. (1998) suggested two test statistics for $\hat{K}$. We attempted both of these tests here, and, except for period 12001–14000, these two tests always lead to the same results. The last column of Table 2 is a summary of the Kaplan test applied to different periods of the return series.[10]

Is the return series i.i.d.? Only for period 16001–18000 is the null of i.i.d. consistently accepted by the three tests. For other series, the null of i.i.d. is either rejected by the ARCH test (six out of ten), or the Kaplan test (six out of ten) or both (three out of ten). Since to show that the EMH is an emergent property, we must have the aggregate result of i.i.d. return series as a precondition, the following analysis of microbehavior will mainly focus on period 16001–18000. However, given the divergence of the test results, the analysis is also extended to all other series as well.

4.2. Mircobehavior

In the previous section, we followed a very standard econometric procedure, and by this procedure, one of the return series is "proved" to be an i.i.d. series, which implies that the price series $\{P_t\}$ from which these return series $\{r_t\}$ are derived is a martingale. So, the EMH in the form of the martingale hypothesis is satisfied in this series. To proceed further and to show that the EMH is an emergent property, one has to ask whether this property can be expected from our understanding of individual behavior. But, the question here is what do we mean by "expected".

One of the advantages of the artificial stock market is that it allows us to observe what traders are actually thinking and doing. For example, we can directly examine whether each trader is effectively a martingale believer. To do so, we simply check the forecasting rules employed by these traders. As argued earlier (Section 2.4), if a trader is effectively a martingale believer, then her forecasting function is simply

$$E_t(P_{t+1} + D_{t+1}) = P_t + E(D_{t+1}) = P_t + \mu, \tag{30}$$

or the function $f_{i,t}$ in Eq. (14) is a zero function. In Section 2.4, we define the number of martingale believers at time t:

$$N_{1,t} = \text{card}\,\{i \,|\, f_{i,t} = 0\}. \tag{31}$$

Therefore, by examining $N_{1,t}$, one can see how well the martingale hypothesis is accepted by traders. Fig. 3 is a time series plot of $N_{1,t}$. From this figure, we can see that most of the time the number of martingale believers is quite small. To have a more precise picture of these numbers, the averages of martingale believers over each subperiods are given in the second column of Table 3. The average corresponding to the period 16001–18000, where the null of i.i.d. is consistently accepted by all tests is only 5.37 (italic-faced in Table 3), i.e. about 1% of the traders. As a result, the martinagle property observed in this period actually comes from a market where few of the participants are martingale believers.

[10] The Kaplan test result is detailed in Appendix C of Chen and Yeh (2000).

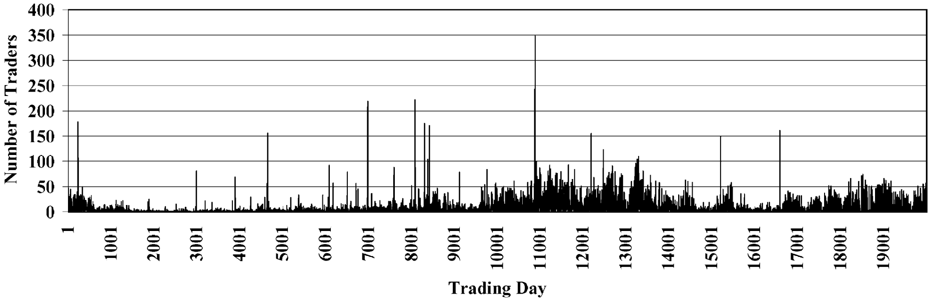

Fig. 3. The number of traders with Martingale strategies on each trading day.

By this inconsistency, we may claim that the martinagle, as an aggregate property, is not expected from our understanding of individual traders. Hence, the EMH is an emergent property.

To make this argument more plausible, and not too much dependent on a single aspect of traders, we also provided evidences in other aspects, namely, traders' search behavior. The efficient market hypothesis has two implications for search. First, there is no need to search, and, second, there is no gain from search. Hence, we do not expect to see too many searchers there, and even if there are, most of them will end up with a futile search. Earlier, we define $N_{2,t}$ as the number of traders who decide to search, and $N_{3,t}$ the number of successful searchers. Here, let the ratio $N_{3,t}/N_{2,t}$ be the proportion of traders who have a successful search, and call it the chance of success, then this chance should be nothing different from throwing a fair coin, which is 0.5. We report the average of $N_{3,t}/N_{2,t}$ over different periods of trading days in the fourth column of Table 3. The rate of success search, on the average, is higher than 50%, and that includes the period 16001-18000 even though its success rate, 51%, is the lowest among the 10 series. Notice also that all these averages are taken over the period with 2000 observations (a quite large sample) and the binomial test

Table 3
Traders' belief and adaptation

Periods	$\bar{N}_1$	$\bar{N}_3$	$\bar{N}_3/\bar{N}_2$	$\bar{k}$	$\bar{\kappa}$
1–2000	3.88	183.71	0.527	7.41	12.57
2001–4000	0.57	183.05	0.528	9.05	15.81
4001–6000	1.37	177.59	0.514	9.03	14.51
6001–8000	2.99	191.98	0.556	9.89	19.13
8001–10000	5.60	184.53	0.535	10.14	22.99
10001–12000	14.02	195.40	0.566	8.61	16.62
12001–14000	12.55	194.06	0.562	8.85	18.23
14001–16000	4.99	177.10	0.514	10.71	24.98
16001–18000	*5.37*	*175.89*	*0.510*	*10.16*	*21.31*
18001–20000	10.16	193.69	0.562	8.72	15.41

The time series is equally decomposed into 10 non-overlapping subperiods. All these numbers are the averages taken over one of these subperiods. $\bar{N}_1$, $\bar{N}_3$, $\bar{N}_3/\bar{N}_2$, $\bar{k}$, and $\bar{\kappa}$ refer to the average number of $N_{1,t}$, $N_{2,t}$, $N_{3,t}/N_{2,t}$, k_t, and κ_t respectively.

(the asymptotic normal test) shows that they are all significantly different from 0.5 though the difference is not greater than 10%. Clearly, search in business school is not entirely futile.

Another way to justify the EMH as an emergent property is to examine the complexity of evolving traders. Intuitively, the EMH says two things about complexity. First, there is no need to be complex. Naive strategies such as martingale ($f_{i,t} = 0$) should be good enough. Second, since there is no need to be complex in the first place, time plays no role, i.e. complexity is not a function of time. To see whether these two observations holds for our traders, we have to first give a definition of complexity, which in general is not an easy job. Fortunately, since all traders' behavior are characterized by their forecasting models, which are in the format of the LISP language, we can easily give two definitions of traders' complexity. 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. For both measures, the complexity degree of the martingale model ($f_{i,t} = 0$) is 1, which can be taken as a benchmark to be compared with the observed behavior of our traders.

On each trading day, we have a profile of the evolved GP-trees for 500 traders, $\{f_{i,t}\}$. Since all forecasting models are in the format of LISP parse trees, their complexity can be measured by the assoicated depth or length. The depth of a parse tree can be defined as the length of the longest path from root to endpoint, whereas the length of a parse tree is measured by counting the number of nodes that appear in the tree (number of elements used in the program). Let $k_{i,t}$ be the depth (the length of the longest path) of the model $f_{i,t}$ and $\kappa_{i,t}$ be the length (the number of nodes) of $f_{i,t}$, then

$$k_t = \frac{\sum_i^{500} k_{i,t}}{500}, \quad \text{and} \quad \kappa_t = \frac{\sum_i^{500} \kappa_{i,t}}{500}. \tag{32}$$

The average of k_t and κ_t are reported in the fifth and sixth column of Table 3. Both figures evidence that traders can evolve toward a higher degree of sophistication, and at some point in time, they can be simple as well. Nevertheless, there is no evidence that traders' behavior will converge to the simple martingale model.

In sum, the three perspectives about traders' behavior, as summarized in Table 3, show that traders are not believers of the EMH, not only in words, but in action as well. While most of the time traders are searching for and using some forecasting models which are much more complex than the simple martingale model, traders do not consider their efforts devoted to these activities futile. As a result, the EMH is anything but what we learned from a analysis of our traders. Like Adam Smith's invisible hand and the Hayek hypothesis, it is an emergent property.

5. The rational expectations hypothesis

The second hypothesis we would like to examine is the rational expectations hypothesis. To some extent, treating the REH as an emergent property seems to be well motivated especially when all traders are given a life which is explicitly boundedly rational. Therefore, it would "surprise" us if these traders can collectively generate a phenomenon that, in a sense, satisfies the REH. However, some qualifications are needed here.

First, the observation that boundedly rational agents can collectively generate phenomena which satisfy the REH is nothing new in economics. Sargent (1993) has an excellent collection of these studies. However, most of these studies were conducted with the device of the representative agent. In this case, they are not suitable for studying the REH as an emergent property. They are a few studies which were conducted without the device of the representative agent and used what is now called agent-based modeling. But, these studies actually showed that the heterogeneity (diversity) of agents eventually disappears and the whole population converges to a representative agent when the aggregate behavior converges to the REH (Arifovic, 1994; Chen and Yeh, 1996). In this case, the REH observed is not an emergent property. Therefore, obtaining the REH as an emergent property is not that trivial.

Second, without the device of the representative agent, it can be quite difficult to define and locate the rational expectations equilibria or fixed points (Spear, 1989; Arthur, 1992; Sargent, 1993), not to mention to observe the REH. Therefore, in this paper, we only give one simple version of the REH, and proceed as follows. We first constructed a representative agent by using the market expectations (objective expectations). The market expectations is defined as the average of all traders' expectations, i.e. $E_t = \sum_{i=1}^{N} E_{i,t}/N$. Given E_t, the prediction error of the representative agent at time t is

$$e_t = E_t - (P_t + E(D_t)). \tag{33}$$

The time series plot of $\{e_t\}$ is given in Fig. 4. We can see that e_t fluctuates around the center 0 in a quite stochastic manner. A browse around the figure shows that the representative agent does not persistently overestimate or underestimate the stock price.

Second, in the spirit of the conventional rational expectations hypothesis test, we assume that the representative agent would not make systematic errors. By systematic errors, we mean that the time series $\{e_t\}$ has patterns. In other words, the time series $\{e_t\}$ is totally unpredictable, or $\{e_t\}$ is an independent series with mean 0. Hence, to test this version of rational expectations hypothesis, we did two things. First, we tested whether the mean forecasting errors are significantly different from 0. The t-statistics are given in the second column of Table 4, which indicate that the null hypothesis is not rejected in all subperiods, and the mean forecasting error is not significantly different from 0. Therefore, the representative agent does not make a systematic error at least in mean (the first moment).

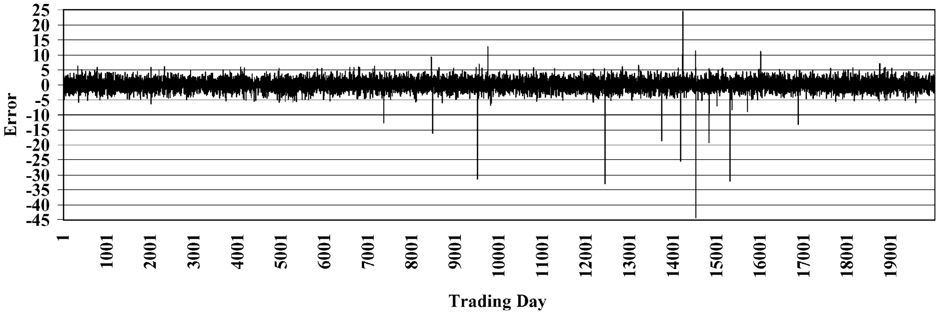

Fig. 4. Time series plot of the forecasting errors of the representative agent.

Table 4
Rational expectations hypothesis and granger causality tests

Period	t-Value	PSC	R^2	BDS	Kaplan	EView	Q-statistic
1–2000	0.03	(2, 0)	0.008	R (11.29)	R	0.00	14.52
2001–4000	0.06	(0, 0)	0.000	A (1.12)	A	0.00	14.54
4001–6000	0.01	(0, 0)	0.000	A (0.85)	A	0.00	14.48
6001–8000	0.07	(1, 0)	0.007	R (2.44)	A	0.00	14.43
8001–10000	0.09	(0, 0)	0.000	A (0.91)	A	0.00	14.37
10001–12000	0.08	(0, 0)	0.000	A (1.63)	A	0.00	13.32
12001–14000	0.07	(0, 0)	0.000	A (1.40)	A	0.00	14.44
14001–16000	0.03	(0, 0)	0.000	A (0.91)	R	0.00	13.38
16001–18000	0.13	(1, 0)	0.006	A (1.41)	A	0.00	14.34
18001–20000	0.06	(0, 0)	0.000	A (1.28)	A	0.00	14.69

The orders p and q are selected based on the Rissanen's PSC criterion. R^2 is the coefficient of determination derived by running the PSC-selected ARMA(p, q) regression. The null hypothesis of the Granger causality test is that D_t does not Granger cause r_t. What is reported in the column of EView is the P-value of the test statistic with 20 lags. The Q-statistic is based on Kau (1997). By the functional central limit theorem and continuous mapping theorem, it can be shown that the 0.05 (0.01) significance level of the Q-statistic is 2.241 (2.807).

Then, one step further, we tested whether the error series is white noise. To do so, we used the PSC filter to extract the linear signal of the error series. The third column of Table 4 gives us the ARMA orders determined by the PSC criterion. From this result, we can see that most of the orders chosen are simply (0, 0). While there are three subperiods (periods 1–2000, 6001–8000, 16001–18000) which are not white noise, from the R^2 given in the next column of the same table, the linear pattern is very weak.

Finally, non-linear patterns. As in Section 5, we used the BDS test and the Kaplan test to test the null of i.i.d. For the BDS test, we only report the result with $\epsilon = 1$ and DIM $= 5$ since it is not sensitive to these two parameters. Based on the BDS test statistics shown in the fifth column of Table 4, the null of i.i.d. is not rejected in most subperiods. As to the Kaplan test, we followed the same procedure introduced in Section 4.1. Two different measures are taken, and if the test results based on these two measures are inconsistent, the decision to reject or accept the null will be ambiguous. The result of the Kaplan test is summarized in the sixth column of Table 4.[11] Like the BDS test, in most series, the null of i.i.d. is not rejected by the Kaplan test.

Putting these four tests together, we find that the null of i.i.d. is not rejected in six out of the 10 series. These six series are bold-faced in Table 4. In other words, there are no systematic patterns of errors made by the representative agent found in these six series; therefore, the REH is not rejected in these six periods.

In sum, what we found is that a collection of boundedly rational agents, through their energetic search and interactions, could in effect generate a representative agent (the "market") who did not make systematic errors in her forecasts. In other words, there is no hidden structure which could be used to improve forecasts but was neglected. In that sense, the REH can be considered another emergent property.

[11] The result of the Kaplan test is detailed in Table 13, Appendix C of Chen and Yeh (2000).

6. Discussion: why emergent?

The main feature of our model that produces the emergent results may be attributed to the use of genetic programming. Unlike GAs, or the specific application of GAs by Arthur et al. (1997), we do not restrict our traders to initialize their search in a small neighborhood of the HREE. This allows us to generate a much larger search space. This larger space can potentially support many forecasting models in capturing short-term predictability, which makes simple beliefs, such as that D_t is an i.i.d. series, or that P_t follows a random walk, difficult to be accepted by traders. In traders' subjective perception, the world is absolutely non-linear, but far from random. This picture has been fully reflected in the complexity statistics (Table 3).

One may suggest that if we slow down the learning speed or reduce the rate of invocation of genetic programming, then the traders will behave more like conventional statisticians who rely heavily on asymptotic theory, and hence will more likely accept those simple beliefs. To test this argument, we ran two experiments. These two experiments were conducted with the same parameters values specified in Table 1, except the evaluation cycle for the business school (m_1) and the evaluation cycle for traders (n_2). For the first experiment, we increased n_2 from 1 to 10, but kept m_1 unchanged; for the second, we kept n_2 at 10 but increased m_1 from 20 to 40. In plain language, in the first experiment we decreased the exploration rate for the traders, but not the b-school, whereas in the latter, we decreased it for both.

Fig. 5 is the time series plot of the evolved complexity of traders' forecasting rules, i.e. depth complexity (k_t) and node complexity (κ_t). By these two experiments, reducing the exploration rate has little effect on the evolved complexity. Instead, the effect of a huge search space on evolved complexity seems to be dominating.

In addition to preventing traders from easily accepting simple beliefs, another consequence of a huge search space is the generation of sunspot-like signals through mutually reinforcing expectations. Traders provided with a huge search space may look for something which is originally irrelevant to price forecasts. However, there is a chance that such kinds of attempts may mutually get reinforced and validated. The generation of sun-spot signals will then drive traders further away from accepting simple beliefs.

To see a version of sunspot-like signals, we take advantage of the assumption that D_t is an i.i.d. series. Earlier, we mentioned that since D_t is assumed to be an i.i.d. series, in the homogeneous rational expectations, it is independent of $\{P_{t+1}\}$, and the history of D_t will not help forecast future price and dividends. However, if a group of traders believe that D_t is not an i.i.d. (contrary to the simple belief), then there is a chance that this initially wrong belief may turn out to be validated by the coordination dynamics of traders. To test whether D_t is a sunspot-like signal in our model, a Granger causality test was applied to the vector time series $\{D_t, r_t\}$ based on our simulated time series shown in Fig. 1.

There are several different ways to conduct the Granger causality test, some tests require an arbitrary choice of filtering processes, and others require an arbitrary choice of lags. In this paper, two versions of the Granger casuality test were applied. One is from the software EView 3.1, which requires an arbitrary choice of lags, and the other is the statistic developed by Kau (1997), which does not require these arbitrary choices.[12]

[12] Kao's test is detailed in Chen and Yeh (2000, p. 26).

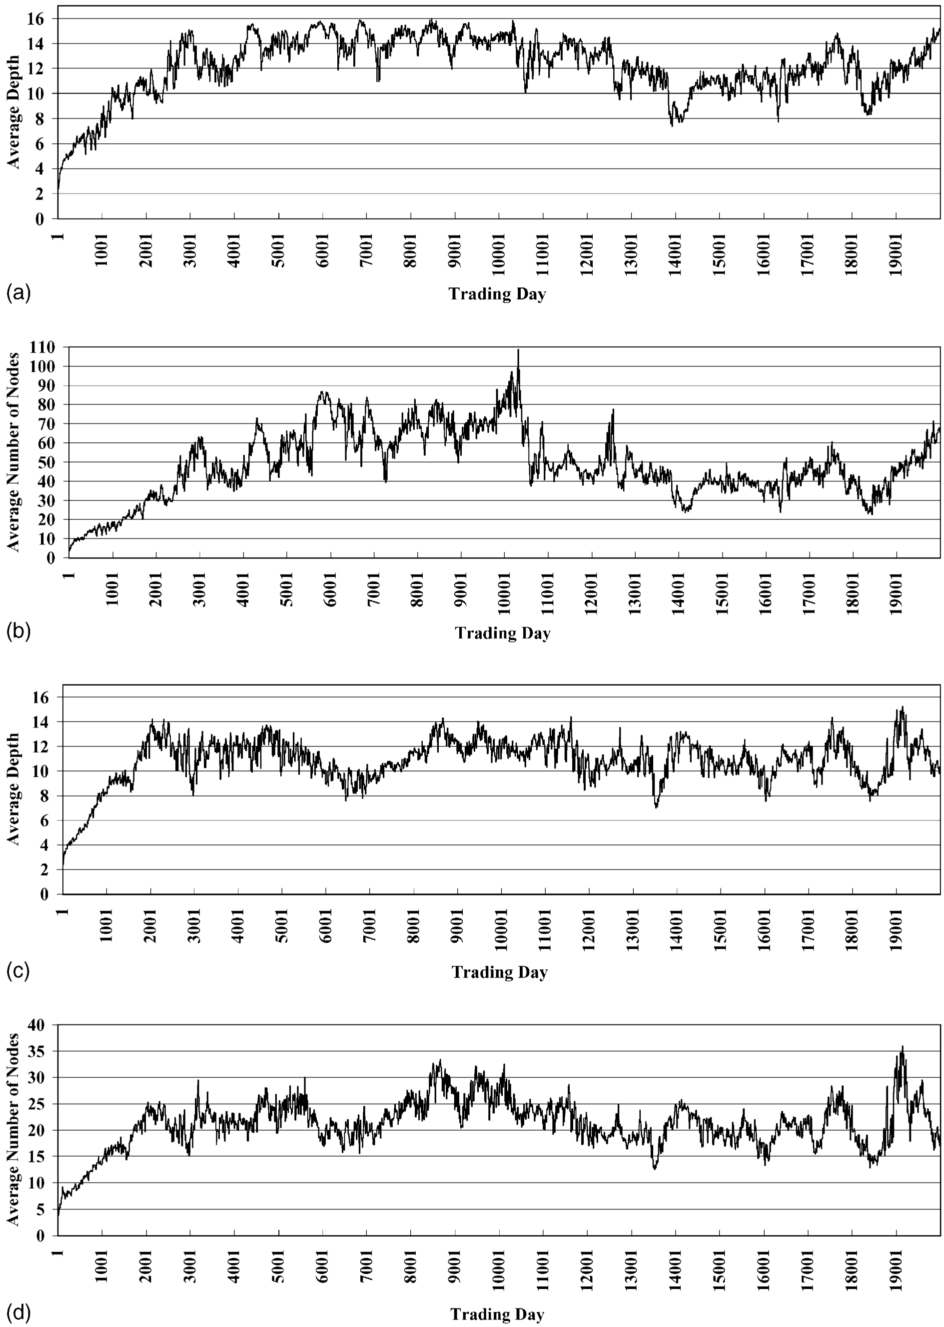

Fig. 5. (a) Depth of complexity (slow learning 1), (b) node of complexity (slow learning 1), (c) depth of complexity (slow learning 2), and (d) node of complexity (slow learning 2).

Table 4 shows the test result of the two statistics mentioned above. The null hypothesis is that D_t fails to Granger cause r_t. From either the P-value or the critical value, we can see that, in all series, both tests consistently reject the null of no causality. Therefore, D_t indeed can help forecast r_t. Since by our experimental design, D_t does not contain the information of future returns, what happened in our simulation is a typical case of mutually supportive expectations that make D_t eventually contain the information of future returns. We believe that sunspot-like signals exist in still other forms which together drive traders further away from simple beliefs.

7. Concluding remarks

By following some standard or modern econometric procedures, this paper examines the aggregate behavior of time series generated by an agent-based artificial stock market. The tests show that some series examined cannot reject a version of the efficient market hypothesis or a version of the rational expectations hypothesis. Thus, we illustrate, to a certain extent, how agent-based models are capable of replicating some well known economic behavior empirically.

However, the properties shown as aggregate results can be quite different from what we observe from the individual behavior. In this paper, the aggregate result of the efficient market can be generated from a collection of interacting traders, most of whom do not believe in the martingale hypothesis (the efficient market hypothesis). Moreover, the aggregate result of the rational expectations can be generated from a collection of boundedly rational agents who behave as if they were never sure about the true model and were continuously searching for a better forecasting model. As a result, the aggregate results are not anticipated from simple scaling-up of the individuals, which is what one may call emergent properties in the literature.

We would like to make a few final remarks about the work done in this study. Firstly, this paper can be read as an extension of the research line stressing that macroeconomic behavior is not a simple scaling-up of microeconomic behavior, e.g. Kirman (1992). More specifically, it is closely related to the agent-based computational model, populated by Epstein and Axtell (1996). This type of model allows us to trace a bottom-up path which is infeasible for conventional models built upon the device of the representative agent, and hence provides an ideal tool to show more precisely how microeconomic behavior can be quite different from macroeconomic behavior. Demonstrating the possibility of this kind of inconsistency is important because it places restrictions on inferring individual behavior from aggregate results. Such restrictions can be critical to policy issues such as decisions to launch a national annuity program based upon a test for the permanent income hypothesis.

Secondly, in terms of agent-based modeling of artificial stock markets, this paper can also be related to the Santa Fe Artificial Stock Market (Arthur et al., 1997). While at this moment this field is too young to define a unified or standard framework, there are still interesting comparisons to be made. First, like the SFI approach, this paper can generate rich varieties of market dynamics. We believe that these types of models offer a promising direction to enrich current studies on microstructure and anomalies. Second, both papers show that the key to understanding the rich dynamics of markets is the mechanism which allows a population of traders to learn or adapt over time. There are lots of parameters which can result in effectively different mechanisms, and it is still in an early stage to evaluate their potential impact.

But, there are also differences between the SFI approach and our approach. The main difference lies in the specific evolutionary computation (EC) technique employed. For them, it is genetic algorithms; for us, it is genetic programming. As Chen and Yeh (1996) asserted, genetic programming can be viewed as a more general way to do the job which used to be done by GAs. However, GP also has its problem when applied to modeling agents' learning. As we have shown in this paper, GP provides us with a large search space and has a great potential to generate sunspot-like signals which can compete with simple beliefs in any finite number of data points through mutually reinforcing dynamics. While we consider this specific design relevant to the emergent properties studied in this paper, to what extent are they empirically relevant is an issue to be pursued in the future.

Acknowledgements

The paper was financially supported by a NSF (Taiwan)-DAAD (Germany) fund. Research grant No. 87-I-2911-15 is gratefully acknowledged. The authors are also grateful to three anonymous referees for their constructive suggestions and to Prof. Thomas Lux for hosting authors' visiting at Bonn University.

References

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

Arthur, W.B., 1992. On Learning and Adaptation in the Economy. Working Paper Series #92-07-038. Santa Fe Institute.

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.

Barnett, W.A., Gallant, A.R., Hinich, M.J., Jungeilges, J.A., Kaplan, D.T., Jensen, M.J., 1998. A single-blind controlled competition among tests for non-linearity and chaos. Journal of Econometrics 82, 157–192.

Brock, W.A., Dechert, W.D., LeBaron, B., Scheinkman, J., 1996. A test for independence based on the correlation dimension. Econometric Reviews 15, 197–235.

Chen, S.-H., Yeh, C.-H., 1996. Genetic programming learning and the cobweb model. Advances in Genetic Programming, Vol. 2, Chapter 22. MIT Press, Cambridge, MA, 1996, pp. 443–466.

Chen, S.-H., Yeh, C.-H., 2001. Evolving traders and the business school with genetic programming: a new architecture of the agent-based artificial stock market. Journal of Economic Dynamics and Control 25, 363–393.

Chen, S.-H., Yeh, C.-H., 2000. On the Emergent Properties of Artificial Stock Markets: The Efficient Market Hypothesis and the Rational Expectations Hypothesis. AI-ECON Research Center Working Paper Series #2000-5. National Chengchi University.

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

Epstein, J., Axtell, R., 1996. Growing Artificial Societies. MIT Press, Cambridge, MA.

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

Hall, R., 1978. Statistical implications of the life cycle-permanent income hypothesis: theory and evidence. Journal of Political Economy 86, 971–988.

Harrald, P., 2000. Phenotype and genotype analogies in evolutionary economic models. In: Chen, S.-H. (Ed.), Evolutionary Computation in Economics and Finance. Physica Verlag (in press).
Holland, J., 1998. Emergence: From Chaos to Order. Addison-Wesley, Reading, MA.
Kaplan, D.T., 1994. Exceptional events as evidence for determinism. Physica D 73, 38–48.
Kau, C.C., 1997. Stock Price–Volume Relationship: New Tests and Their Applications, Master Thesis. Department of Economics, National Taiwan University.
Kirman, A., 1992. Whom or what does the representative individual represent. Journal of Economic Perspectives 6 (2), 117–136.
Koza, J., 1992. Genetic Programming: On the Programming of Computers by Means of Nature Selection. MIT Press, Cambridge, MA.
Krugman, P., 1996. The Self-Organizing Economy. Blackwell (Basil), Oxford.
LeBaron, B., 2000. Agent-based computational finance: suggested readings and early research. Journal of Economic Dynamics and Control 24 (5–7), 679–702.
LeBaron, B., Arthur, W.B., Palmer, R., 1999. Time series properties of an artificial stock market. Journal of Economic Dynamics and Control 23, 1487–1516.
Lux, T., 1995. Herd behavior, bubbles and crashes. Economic Journal 105 (431), 881–896.
Lux, T., 1997. Time variation of second moments from a noise trader/infection model. Journal of Economic Dynamics and Control 22, 1–38.
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, 498–500.
Mayr, E., 1997. This is Biology: The Science of the Living World. Harvard University Press, Cambridge, MA.
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.
Sargent, T.J., 1993. Bounded Rationality in Macroeconomics. Oxford University Press, Oxford.
Spear, S.E., 1989. Learning rational expectations under computability constraints. Econometrica 57 (4), 889–910.
Tversky, A., Kahneman, D., 1974. Judgment under uncertainty: heuristics and biases. Science 185, 1124–1131.