Simulating the Evolution of Portfolio Behavior in a Multiple-Asset Agent-Based Artificial Stock Market

Shu-Heng Chen[1] and Ya-Chi Huang[1]

AI-ECON Research Center, Department of Economics,
National Chengchi University, Taipei, Taiwan 11623
chchen@nccu.edu.tw, yachi@aiecon.org
http://aiecon.org/

Abstract. [Blume and Easley (1992)] show that if agents have the same savings rule, an expected discounted logarithmic utility maximizer with correct beliefs will dominate. If no agent adopts this rule, then agents with incorrect beliefs, but equally averse to risk as logarithmic utility maximizers, may eventually hold more wealth than the agent with correct beliefs. In other words, a trader with correct beliefs can be driven out of the market by traders with incorrect beliefs.
However, [Sandroni (2000)] shows that, among agents who have the same intertemporal discount factor and who choose savings endogenously, the most prosperous will be those making accurate predictions. Agents with inaccurate predictions will be driven out of the market regardless of their preferences.
By using the extended agent-based artificial stock market, we simulate the evolution of portfolio behavior, and investigate the characteristics of the long-run surviving population of investors.
Our agent-based simulation results are largely consistent with [Blume and Easley (1992)], and we conclude that preference is the key factor determining agents' survivability.

1 Motivation

A long-standing theory in economics is that agents who do not predict as accurately as others are driven out of the market, and it underlies the efficient-markets hypothesis and the use of rational expectations equilibrium as a solution concept because it implies that asset prices will eventually reflect the beliefs of agents making accurate predictions.

The recent literature casts serious doubt on the theory that agents with incorrect beliefs will be driven out of the market by those with correct beliefs. [Blume and Easley (1992)] is an example.

[Blume and Easley (1992)] first showed that the only portfolio rule that can survive in the long run is the one which maximizes the expected growth rate of wealth share accumulation. Based on the Kelly criterion ([Kelly (1956)]), this type of behavior is equivalent to that of maximizing a logarithmic utility function. In other words, the fittest investor should behave as if he were endowed with a logarithmic utility function. The rule is also called the maximizing expected logarithm utility rule, or, briefly, the MEL rule. Blume and Easley showed that if there exists an investor who uses the MEL rule in the market, then other rational investors will not survive. If no agent adopts this rule, then agents with incorrect beliefs, but as equally averse to risk as the logarithmic utility maximizer, may eventually hold more wealth than the agent with correct beliefs. So, even though an investor has perfect foresight, he may not survive simply because his utility function is not logarithmic.

Blume and Easley's finding is surprising because few economic studies have ever regarded the preference as a primary force in determining an agent's survivability, and have downplayed the role of beliefs. They obliquely controvert the efficient-markets hypothesis.

The Blume-Easley result was questioned by [Sandroni (2000)]. He considered the portfolio rule to be only a part of the investment decision, which is not complete if savings are left out (exogenously fixed). He developed a more general framework in which both savings and the portfolio rule were derived from maximizing expected discounted utility, and showed that the most prosperous will be those who make accurate predictions. Agents with inaccurate predictions will be driven out of the market, and therefore convergence to rational expectations is obtained. The surviving agents may have diverse preference over risk. In other words, the utility function does not play a role in determining survivability.

However, despite the title of the article, the approach taken by [Blume and Easley (1992)] is not really evolutionary in the sense that finite-time survival pressure is not put on investors ([Farmer and Lo (1999)], [Tesfatsion (2001)]). Instead, their results are built upon the *asymptotic analysis* in probability. This approach is mathematically rigorous, but has severe restrictions. The results derived from the asymptotic analysis may fail to predict what would actually happen in an evolutionary process. Recent studies, such as [Lettau (1997)], and [LeBaron (2001)], have shown that the market dynamics can behave nontrivially differently and are path-dependent when finite-time survival pressure is placed on agents. Second, the role of *learning* and *adaptation* in this debate is largely neglected in [Blume and Easley (1992)] and [Sandroni (2000)].[1] Introducing learning and adaptation into Sandroni's model may make the resultant stochastic process so complex that few analytical results can be obtained.

Furthermore, both [Blume and Easley (1992)] and [Sandroni (2000)] are analytical papers with various technical conditions that are hard to verify in the real world. For example, [Blume and Easley (1992)] supported the Kelly criterion, which states that a fit rule should maximize the expected growth rate of wealth and that the rule of an expected logarithmic utility maximizer is one

[1] In their Section 6 "Adaptive Behavior," Blume and Easley did extend their analysis to a few types of adaptive behavior. However, that was done with the assumption that *all investors are endowed with logarithmic utility functions*, and hence they are not very helpful.

simple class of this kind of rule. However, the Kelly criterion is obtained under specific assumptions about the price process which would be hard to verify in an equilibrium model where the price process is endogenously determined. Furthermore, Sandroni's analysis assumes that all utility functions satisfy the Inada condition. Unfortunately, this assumption excludes some interesting preferences frequently employed in financial economics.

These limitations render the empirical parts of the papers empty:

- First, does risk attitude (risk preference) matter?
- Second, does forecasting accuracy matter?
- Third, would wealth distribution tend to degenerate to certain groups of people or remain randomly distributed?
- Fourth, what happens in the case of other familiar types of traders, such as CAPM traders?

The purpose of our paper is to use a computational model, or more precisely, the agent-based computational model, to verify what has been said on this subject. Agent-based modeling places great emphasis on the cognitive limits of traders and their associated learning processes. It allows us a real-time (and finite time, not time in the limit) evaluation of the prophecies made above. The path-dependence issue, usually due to the learning and search process, is encapsulated into this model. Therefore, we believe that a useful contribution to the debate can come from the adoption of agent-based computational modeling ([Arthur et. al. (1997)], [LeBaron (2000)]). Agent-based computational modeling of financial markets allows us to extend Sandroni's model to a market consisting of boundedly-rational heterogeneous investors (heterogeneous in both beliefs and preferences). Being described as "...a truly new frontier whose exploration has just begun" by [Farmer and Lo (1999)], the agent-based artificial stock market is far from its fully-fedged state. To date, all studies have concentrated on the single-asset version of it. The multi-asset version is still not available. Therefore, this study, to the best of our knowledge, may be the first attempt to make such progress in this direction.

2 The Model

Consider a complete securities market. Time is discrete and indexed by $t = 0, 1, 2, ...$ There are M states of the world indexed by $m = 1, 2, ..., M$, one of which will occur at each date. States follow a stochastic process, which is characterized by the first-order Markov process or iid process. Asset $m \in \{1, 2, ...M\}$ pays $w_m > 0$ when state $m \in \{1, 2, ...M\}$ occurs, and 0 otherwise. At each date t, there is only one unit of each asset available, so that the total wealth in the economy at date t will be $W_t = w_m$, if state m occurs. The wealth will be distributed among the investors proportionately according to their owned share of asset m. The distribution received by each investor can be used to consume and re-invest. We assume that there is aggregate uncertainty so that $w_m \neq w_v$, for $m \neq v$. Let $\rho_{m,t}$ be the market price of asset m at date t.

There is a finite number of traders with *heterogeneous* preferences in this economy, indexed by $i \in \{1, 2, ...I\}$. Each trader i has J beliefs over M states at each date t, with the beliefs being denoted by $\{B^i_{j,t}\}^J_{j=1}$, and he will choose the best at each date t. Every Δ periods, each agent will update his belief set via a *High Level GA*. According to his best belief j, the investor i will determine his optimal saving rate $\delta^{i,*}_{j,t}$ and portfolio weights $\alpha^{i,*}_{j,t} = (\alpha^{i,*}_{j,1,t}, \alpha^{i,*}_{j,2,t}, ..., \alpha^{i,*}_{j,M,t})$ by maximizing his lifetime expected discounted utility via a *Low Level GA*. We denote investor i's wealth at time t by W^i_t. Given his saving rate $\delta^{i,*}_{j,t}$, the agent will, therefore, invest a total of $\delta^{i,*}_{j,t} \cdot W^i_{t-1}$ in the M assets at time t. Furthermore, if we let $q^i_{m,t}$ be agent i's demand for (shares of) asset m at time t, then

$$q^i_{m,t} = \frac{\alpha^{i,*}_{j,m,t} \cdot \delta^{i,*}_{j,t} \cdot W^i_{t-1}}{\rho_{m,t}}, \quad m = 1, 2, ..., M \tag{1}$$

To be more precise regarding the terminology, we shall refer to $\alpha^{i,*}_{j,t}$ as trader i's optimal ***portfolio rule*** (under belief j at date t) and the pair $\{\delta^{i,*}_{j,t}, \alpha^{i,*}_{j,t}\}$ as trader i's optimal ***investment rule***. We shall now omit the notation j as long as the effect of the belief on investment is clear from the context.

In equilibrium, prices must be such that markets clear, i.e. total demand equals total supply:

$$\sum_{i=1}^{I} \frac{\alpha^{i,*}_{m,t} \cdot \delta^{i,*}_{t} \cdot W^i_{t-1}}{\rho_{m,t}} = 1, \quad m = 1, 2, ..., M \tag{2}$$

Therefore, the market equilibrium price of asset m will be determined by

$$\rho_{m,t} = \sum_{i=1}^{I} \alpha^{i,*}_{m,t} \cdot \delta^{i,*}_{t} \cdot W^i_{t-1}, \tag{3}$$

and agents' shares of assets will be determined accordingly by Equation (1). The actual state at date t will realize, and the wealth will be distributed among the investors proportionately according to their owned share of asset m if state m occurs. Therefore, agent i's wealth at date t will be determined by $W^i_t = q^i_{m,t} \cdot w_m$.

3 Research Strategies and their Implementation

In this research project, we shall first *extend* the current *single-asset* artificial stock market to its multiple-asset version in which we can simulate the evolution of portfolio behavior. The analytical models upon which our artificial stock market is built are those of [Blume and Easley (1992)] and [Sandroni (2000)]. In a way, this research project can be regarded as an agent-based version of [Sandroni (2000)].

3.1 Agents' Cognition

Like all agent-based computational economic models, we shall first start with a description of a *typical* agent, including his *cognition* and *adaptive behavior*. Let us first start with the problem presented to our agents. Agents in our model behave like a normal investor who tries to maximize his lifetime discounted expected utility by appropriately choosing his *investment strategy*. The investment strategy is mainly composed of two parts, namely, *saving* and *portfolio*.

At each point in time, say t, the investor i observes a time series (history) of the realization of the states, namely, $S_{t-1} \equiv \{m_s\}_{s=0}^{t-1}$ ($m_s \in \{1, 2, ..., M\}$). Based on this realization S_{t-1}, he makes his decisions on a sequence of investment strategies:

$$\{\{\delta^i_{t+r}\}^{\infty}_{r=0}, \{\alpha^i_{t+r}\}^{\infty}_{r=0}\},$$

where δ^i_t is the saving rate at time t, and α^i_t is the portfolio comprising the M assets. Given investor i's temporal utility function u^i, it is hoped that this sequence of investment strategies is rational in the sense that his lifetime discounted expected utility can be maximized (see Equation 4). Mathematically, the optimization problem can be stated as follows:

$$\max_{\{\{\delta^i_{t+r}\}^{\infty}_{r=0}, \{\alpha^i_{t+r}\}^{\infty}_{r=0}\}} E\{\sum_{r=0}^{\infty} (\beta^i)^r u^i(c^i_{t+r}) \mid S_{t-1}\} \tag{4}$$

subject to

$$c^i_{t+r} + \sum_{m=1}^{M} \alpha^{i,*}_{m,t+r} \cdot \delta^{i,*}_{t+r} \cdot W^i_{t+r-1} \leq W^i_{t+r-1} \quad \forall r \geq 0, \tag{5}$$

$$\sum_{m=1}^{M} \alpha^i_{m,t+r} = 1, \quad \alpha^i_{m,t+r} \geq 0 \quad \forall r \geq 0. \tag{6}$$

The c^i_t is the consumption of investor i at time t, which satisfies Equation (7).

$$c^i_{t+r} = (1 - \delta^i_{t+r}) W^i_{t+r-1} \tag{7}$$

Constraint (5) is simply the borrowing constraint: agents in our model cannot invest by borrowing.

Solving Equation (4) requires that investors have no cognitive limit. For example, they have perfect foresight so that the future prices of asset m, $\{\rho_{m,t+r}\}_{r=0}^{\infty}$ are known to them. In addition, they can correctly infer from the past realization S_{t-1} the stochastic process which generated $\{m_t\}$. Unfortunately, none of these can be satisfied in the real situation. Moreover, in markets composed of complex heterogeneous agents, the rational expectations equilibria may not even be computable ([Spear (1989)], [Tay and Linn (2001)]). This provides room for other approaches, such as *adaptive computing*.

Over the last few years, the *genetic algorithm* has been the most active tool in adaptive computing. In agent-based computational economics, it is mainly used to deal with either the cognitive limit of *optimizing*, or the cognitive limit of *forecasting*. Very few studies use the GA to conduct *multi-level* evolution, and this research project purports to be a pioneering attempt in this direction: we use the genetic algorithm to evolve agents' *investment strategies* and *beliefs* simultaneously. The two-level evolution proceeds as follows:

– At a fixed time horizon, investors update (evolve) their beliefs of the states coming in the future.
– They then evolve their investment strategies based on their beliefs.

The two-level evolution allows agents to solve a *boundedly-rational* version of the optimization problem (4). First, the cognitive limit of investors and the resultant adaptive behavior free them from an infinite-horizon stochastic optimization problem, as in Equation (4). Instead, due to their limited perception of the future, the problem effectively posed to them is the following:

$$\max_{\{\{\delta_{t+h}\}_{h=0}^{H-1},\{\alpha_{t+h}\}_{h=0}^{H-1}\}} E\{\sum_{h=0}^{H-1}(\beta^i)^h u^i(c^i_{t+h}) \mid B^i_t\} \tag{8}$$

Here, we replace the infinite-horizon perception with a finite-horizon perception of length H, and the filtration (σ-algebra) induced by S_{t-1} with B^i_t, where B^i_t is investor i's *belief* at date t. In a simple case where m_t is *independent* (but not necessarily stationary), and this is known to the investor, then B^i_t can be just the *subject probability function*, i.e.

$$B^i_t = (b^i_{1,t}, ...b^i_{M,t}), \tag{9}$$

where $b^i_{m,t}$ is investor i's subjective probability of the occurrence of the state m in any of the next H periods. The expectation of $u^i(c^i_{t+h})$, for all h, will then be taken with respect to the probability function (9). In a more general setting, B^i_t can be a *high-order Markov process*. With this replacement, we assume that investors have only a vague perception of the future, but will continuously adapt when approaching it. As we shall see in the second level of evolution, B^i_t is *adaptive*.

Furthermore, we assume that investors will *continuously* adapt their investment strategies according to the *sliding window* shown in Figure 1. At each point

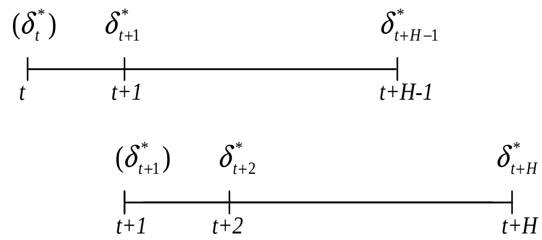

Fig. 1. A Sliding-Window Perception of the Investors

in time, the investor has a perception of a time horizon of length H. All his investment strategies are evaluated within this reference period. He then makes his decision based on what he considers to be the best strategy. While the plan comes out and covers the next H periods, only the first period, $\{\delta_t^{i,*}, \alpha_t^{i,*}\}$, will be actually implemented. The next period, $\{\delta_{t+1}^{i,*}, \alpha_{t+1}^{i,*}\}$, may not be implemented because it may no longer be the best plan when the investor receives the new information and revises his beliefs.

With this sliding-window adaptation scheme, one can have two further simplifications of the optimization problem (4) – (6). The first one is that the future price of the asset m, $\rho_{m,t+h}$ remains unchanged for each experimentation horizon, namely, at time t,

$$\rho_{m,t+h}^i = \rho_{m,t-1}, \quad \forall \quad h \in \{0, H-1\}, \tag{10}$$

where $\rho_{m,t+h}^i$ is i's subjective perception of the h-step-ahead price of asset m. Second, the investment strategies to be evaluated are also time-invariant under each experimentation horizon, i,e.

$$\delta_t^i = \delta_{t+1}^i = \delta_{t+2}^i = ...\delta_{t+H-1}^i, \tag{11}$$

$$\alpha_t^i = \alpha_{t+1}^i = \alpha_{t+2}^i = ...\alpha_{t+H-1}^i. \tag{12}$$

With these two simplifications, we replace the original optimization problem, (4) – (6), that is presented to the infinitely-smart investor, with a modified version which is suitable for a boundedly-rational investor.

$$\max_{\{\{\delta_t^i\},\{\alpha_t^i\}\}} E\{\sum_{h=0}^{H-1} (\beta^i)^h u^i(c_{t+h}^i) \mid B_t^i\} \tag{13}$$

subject to

$$c_{t+h}^i + \sum_{m=1}^{M} \alpha_{m,t+h}^{i,*} \cdot \delta_{t+h}^{i,*} \cdot W_{t+h-1}^i \leq W_{t+h-1}^i, \quad \forall \quad h \in \{0, H-1\}, \tag{14}$$

$$\sum_{m=1}^{M} \alpha^i_{m,t} = 1, \alpha^i_{m,t} > 0, \forall m, \quad (15)$$

$$c^i_{t+h} = (1-\delta^i_t)W^i_{t+h-1}, \ \ \forall \ \ h \in \{0, H-1\}. \quad (16)$$

3.2 Evolution at the Low Level: Investment Strategies

Coding and Initialization A solution for the finite-horizon stochastic optimization problem, (13) – (16), for all well-behaved utility functions, can be effectively and automatically solved with the *genetic algorithm*. The implementation of the genetic algorithm starts with a representation (coding) of solutions. Here, we employ the real coding (the direct coding). The saving rate (δ^i_t) and the portfolio (α^i_t) will not be coded as bit strings, but as real-valued numbers:

$$\{\delta^i_t \mid \alpha^i_{1,t}, \alpha^i_{2,t}, ..., \alpha^i_{M,t}\} \quad (17)$$

To solve (13), an initial population of investment strategies with *population size* N is first generated for each investor i,

$$GEN^i_{t,0} \equiv \{\delta^i_{t,n}(0), \alpha^i_{t,n}(0)\}^N_{n=1}.$$

The number inside the parentheses refers to the generation number in the GA cycle. Population $GEN^i_{t,0}$ is generated as follows:

- $\delta^i_{t,n}(0)$ is randomly generated from the uniform distribution $U(0,1)$.
- To generate a portfolio $\alpha^i_{t,n}(0)$, a set of numbers

  $$(Q_1, Q_2, ..., Q_M)$$

  are randomly generated from $U(0,1)$. Then, to make sure that their sum is equal to 1, they are rescaled as follows:

  $$(\frac{Q_1}{\sum^M_{q=1} Q_q}, \frac{Q_2}{\sum^M_{q=1} Q_q}, ..., \frac{Q_M}{\sum^M_{q=1} Q_q}) \quad (18)$$

Fitness Evaluation: Eval { $\boldsymbol{GEN^i_{t,g}}$ **}** Corresponding to (13), the fitness measure f is simply the H-horizon discounted expected utility:

$$f_t(n,g) \equiv f(\delta^i_{t,n}(g), \alpha^i_{t,n}(g)) \equiv E\{\sum_{h=0}^{H-1} (\beta^i)^h u^i(c^i_{t+h}) \mid B^i_t\}, \quad (19)$$

where $f_t(n,g)$ refers to the fitness of the nth investment strategy in the population $GEN^i_{t,g}$ (i.e. the gth generation of the GA cycle). The Monte Carlo simulation technique is used to evaluate the fitness (19). The way to do so is to simulate a large number, say L, of H-horizon histories of the states based on investor i's belief, B_t^i. For each *simulated history* l ($l \in [1, L]$), we can obtain a realization of (19), i.e.

$$\sum_{h=0}^{H-1} (\beta^i)^h u^i(c_{t+h}^i \mid l), \quad l = 1, 2, ...L.$$

Then, estimating $f_t(n, g)$ involves taking the sample average,

$$\hat{f}_t(n, g) = \frac{\sum_{l=1}^{L} \sum_{h=0}^{H-1} (\beta^i)^h U^i(c_{t+h}^i \mid l)}{L}. \tag{20}$$

Genetic Operation: $GEN_{t,g}^i \rightarrow GEN_{t,g+1}^i$ Once the procedure **Eval {** $GEN_{t,g}^i$ **}** is completed, all investment strategies are associated with a fitness which is the output of (20).

$$\textbf{Eval}: \ \{\delta_{t,n}^i(g), \alpha_{t,n}^i(g)\}_{n=1}^N \rightarrow \{f_t(n, g)\}_{n=1}^N \tag{21}$$

Based on their fitness, we shall revise and renew these investment strategies based on investor i's belief B_t^i. This revision and renewal procedure involves the use of four standard genetic operators, namely, *selection*, *crossover*, *mutation* and *election*.

Selection: The *tournament selection* with tournament size 4 is employed. For each selection, four investment strategies are randomly selected from $GEN_{t,g}^i$. Of these, the best two will be chosen as the parents (mating pool). We denote them by

$$I_x \equiv \{\delta_{t,x}^i(g), \alpha_{t,x}^i(g)\},$$

and

$$I_y \equiv \{\delta_{t,y}^i(g), \alpha_{t,y}^i(g)\},$$

where $x, y \in [1, N]$.

Crossover: With probability one (*crossover rate =1*), the two parents chosen above will generate an offspring by taking a weighted average of the two investment strategies, and the weights will be determined by the relative fitness of the two strategies.

$$\begin{aligned} I_z &\equiv (\delta_{t,z}^i(g), \alpha_{t,z}^i(g)) \\ &= \frac{f_t(x,g)}{f_t(x,g)+f_t(y,g)}(\delta_{t,x}^i(g), \alpha_{t,x}^i(g)) + \frac{f_t(y,g)}{f_t(x,g)+f_t(y,g)}(\delta_{t,y}^i(g), \alpha_{t,y}^i(g)) \end{aligned} \tag{22}$$

Mutation: The offspring I_z will then have a small probability (mutation rate) to mutate. If mutation happens, it will proceed as follows. For the saving rate, a number randomly selected from the $U[0, 1]$ will be used to replace $\delta_{t,z}^i(g)$. For the portfolio, a set of numbers,

$$\epsilon \equiv (\epsilon_1, \epsilon_2, ..., \epsilon_M),$$

randomly generated from $U(0,1)$, will be added to $\alpha^i_{t,z}(g)$ component by component. Then the rescaling technique described in (18) will be applied. We call the resultant strategy $I_{z'}$.

Election: The use of the election operator guarantees that the new investment strategy is expected to perform better than the ones it replaced. In election, we shall use (20) to evaluate the potential fitness of $I_{z'}$, and compare it with the fitness of the two parents, I_x and I_y. Then, only the one with the highest fitness will be retained for the next generation, $GEN^i_{t,g+1}$.

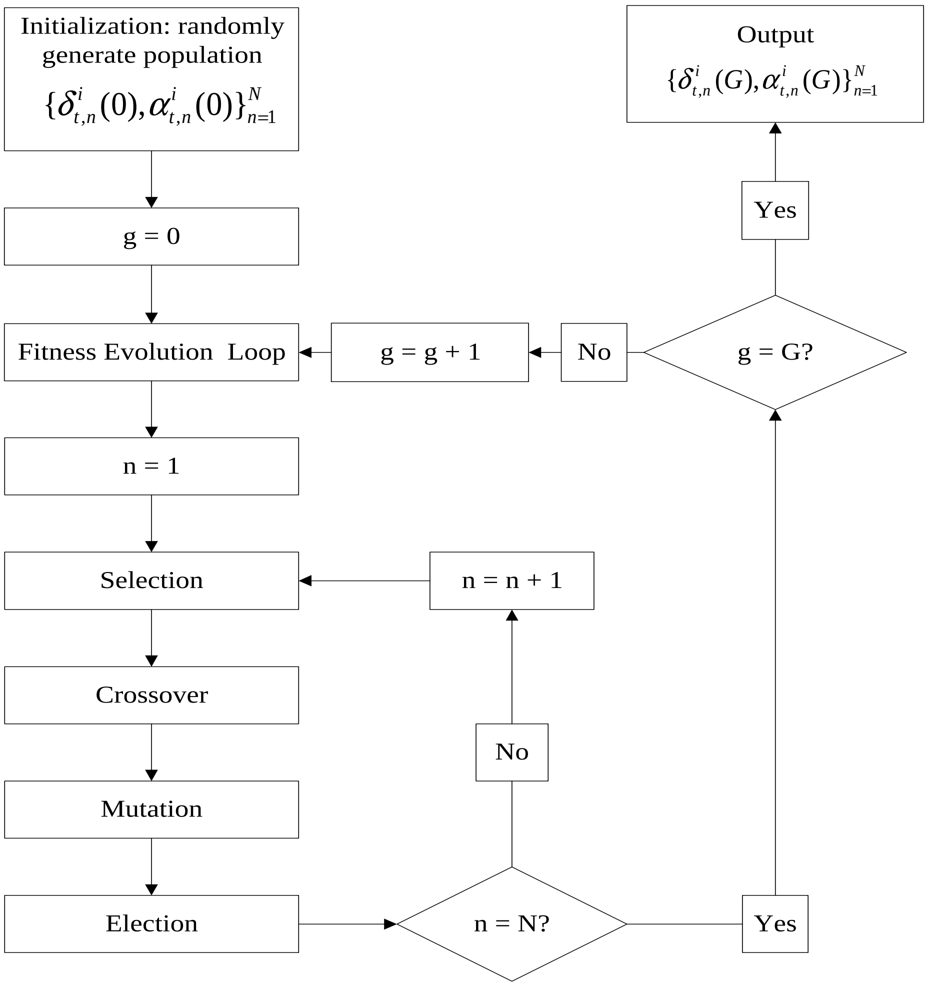

Fig. 2. Flowchart of the Low-Level GA

Loops Once a new investment strategy is generated, a *loop* (Figure 2) leads us back to selection, which is then followed by crossover, mutation and election and then the next new investment strategy is generated. The loop will continue until all N strategies of $GEN^i_{t,g+1}$ are generated. $GEN^i_{t,g+1}$ will be evaluated based on the **Eval** procedure, and based on the evaluation, genetic operators will be applied to $GEN^i_{t,g+1}$ to generate $GEN^i_{t,g+2}$. This loop will also be repeated over and over again until a termination criterion is met, e.g., when g reaches a prespecified number G.

When the renewal and revision process is over, the investor will select the best strategy from the last population of investment strategies, say, $GEN^i_{t,G}$.

$$(\delta_t^{i,*}, \alpha_t^{i,*}) = \arg \max_{GEN^i_{t,G}} \{f_t(n, G)\}_{n=1}^N \tag{23}$$

Except for the *CAPM believers*, the procedure described above, as summarized in Figure 3, will be applied to all investors to generate their optimal investment decisions. The behavior of CAPM believers will be discussed later.

3.3 Evolution at the High Level: Beliefs

At the low level of evolution, the investor revises and renews his investment strategies with respect to a specific belief selected from *a population of beliefs* $\{B^i_{j,t}\}_{j=1}^J$. In other words, at each point in time, the investor may have more than one model of uncertainty in the world. The idea that each agent can simultaneously have several different models of the world, which are competing with each other in a co-evolving process, is a distinguishing feature of the *population learning models* ([Holland and Miller (1991)], [Arthur et. al. (1997)], [Vriend (2000)]). Of course, these models are not equally promising, and the investor tends to base his decision (investment strategies) on one of the most promising ones. However, as times goes on, his beliefs of the world will be revised and renewed in light of the newly incoming information. In this section, we shall describe how genetic algorithms can be applied to modeling the beliefs updating process.

Coding and Initialization In the Blume-Easley-Sandroni model, each investor's perception of the uncertainty (finite-state stochastic process) of the market can be characterized by two elements: first, the *dependence structure* (k), and, second, the *sample size* (d). Based on this characterization, the investor believes that the market over the last d days follows a kth-order Markov process. According to this belief, he would use $\{m_{t-s}\}_{s=v+1}^{v+d+1}$ to estimate the Markov transition matrix where the notation v will be introduced later. As a result, each belief can be represented by a binary string, of length $\tau_1 + \tau_2$,

$$\underbrace{a_1 a_2 ... a_{\tau_1}}_{\tau_1 bits} \underbrace{a_{\tau_1+1} a_{\tau_1+2} ... a_{\tau_1+\tau_2}}_{\tau_2 bits}, \quad a_i \in \{0, 1\}, \quad \forall 1 \le i \le \tau_1 + \tau_2$$

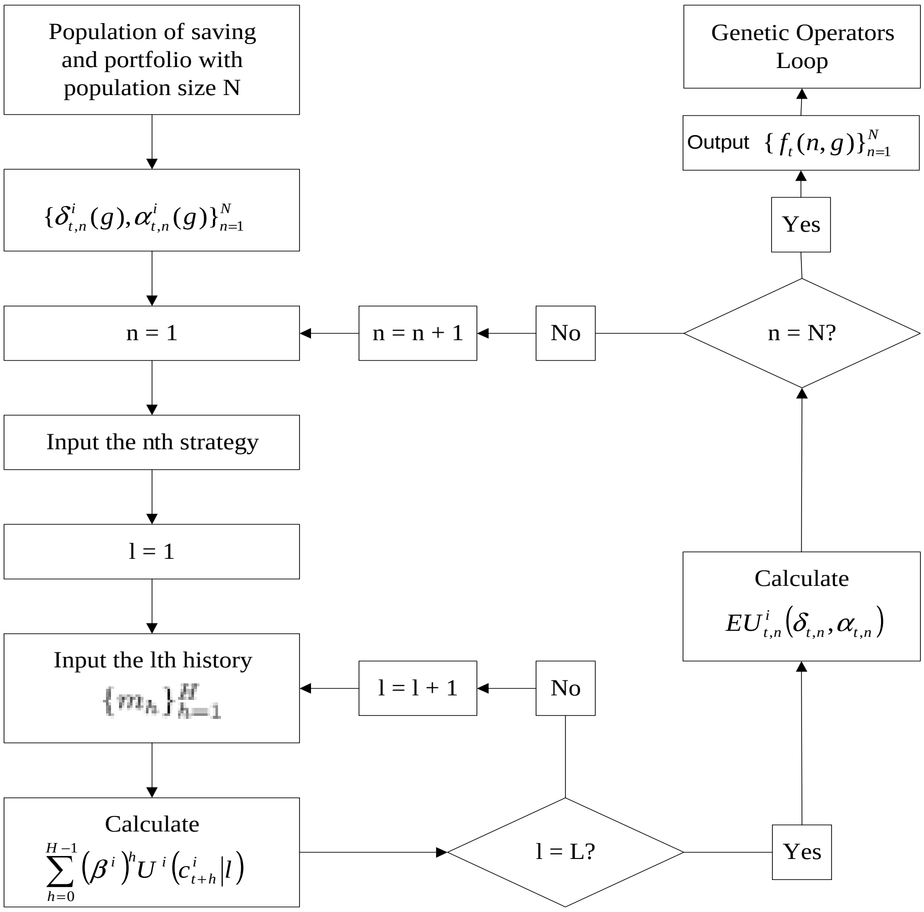

Fig. 3. The Flowchart of the Investment Optimization

that has the following interpretation: the states follow a Markov process of the order

$$k = (\sum_{i=1}^{\tau_1} 2^{\tau_1 - i} a_i) \tag{24}$$

over the last

$$d = (\sum_{i=\tau_1+1}^{\tau_1+\tau_2} 2^{\tau_1+\tau_2-i} a_i) + c \tag{25}$$

days in what we have referred to as the training period which will be specified later, where c is the minimum number of observations required for the estimation. In our current model, we simplify and limit the *dependent structure (k)* to 0 or 1 , that is, we only assume the stochastic process to be iid or Markov. Therefore, we assume that $c = 10$ and $\tau_1 = 1$, $\tau_2 = 9$ throughout the paper.

At the initial date ($t = 0$), all investors are endowed with a population of J beliefs (J bits), which are randomly generated. Out of these J beliefs, the investor will randomly choose one on which his investment decision is based. Then every Δ days, this population of belief will be reviewed and revised based on the fitness function.

Belief Updating Scheme Agents in our model behave like machine-learning agents. They are assumed to know the risk of over-fitting, and hence use validated data to perform the model selection. One way of ensuring that our agents behave so is to set the fitness function as the fitting error in the validation set, rather than the training set. Therefore, we design the belief updating scheme as in Figure 4.

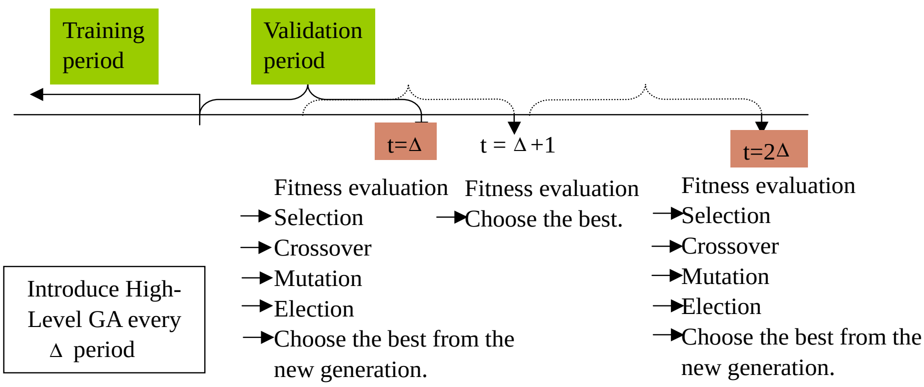

Fig. 4. The Belief Updating Scheme

As we can see from this figure, at each time t agents retain the most recent v days as the validation period. They use the data before the validation period, that is, the data of the training period, to estimate the parameters of each belief. Then a fitness measure for the probability function is the *likelihood*,

$$L^i_{j,t} = L(\{m_{t-s}\}^v_{s=1} \mid B^i_{j,t}), \tag{26}$$

where $\{m_{t-s}\}^v_{s=1}$ is the state history over the last v days. Equation (26) is the likelihood of the observations $\{m_{t-s}\}^v_{s=1}$ in the validation period under the belief $B^i_{j,t}$. Every Δ period, after they finish the evaluation of each belief's fitness, they apply the genetic operation to update their belief set, and the belief with the highest fitness will be chosen. Even in the period that the genetic operation is not applied, they evaluate the fitness of beliefs in their current belief set using the newest data and choose the best from it.

Genetic Operation Once the procedure **Eval** $\{B^i_{j,t}\}^J_{j=1}$ is completed, all beliefs are associated with a fitness which is the output of (26).

$$\textbf{Eval} : \{B^i_{j,t}\}^J_{j=1} \rightarrow \{L^i_{j,t}\}^J_{j=1} \tag{27}$$

Based on this fitness evaluation, we will revise and renew investor i's beliefs by using the following four genetic operators: selection, crossover, mutation and election.

Selection: A tournament selection with tournament size 4 is adopted. The best two beliefs will be chosen as the parents (mating pool).

Crossover: With probability one ($crossoverrate = 1$), the two parents chosen above will generate an offspring by the *uniform crossover*. With this crossover, each bit position of the offspring will be taken randomly either from the father or the mother with a one-half chance for each. For an illustration, let us consider the pair of parents to be $B^i_{x,t} = 0010101010$ and $B^i_{y,t} = 0111110010$. Then, an offspring, B^i_z, can be

$$B^i_z = 0011100010 \rightarrow (k_z, d_z) = (0, 240).$$

Mutation: There is a small probability (mutation rate) by which each bit of B^i_z may encounter a change. For example, the mutation which changes the first bit from "0" to "1", and the fifth bit from "1" to "0" will result in a new string:

$$B^i_{z'} = 0011000010 \rightarrow (k_{z'}, d_{z'}) = (0, 208).$$

Election: Finally, $B^i_{z'}$ will also be evaluated by the observations $\{m_{t-s}\}^v_{s=1}$, and the likelihood will be figured out. We will then compare the likelihood from $B^i_{z'}$ with the likelihood from the parent models, and the best one will be passed to the next generation, $\{B^i_{j,t}\}^J_{j=1}$.

Loops Once a belief is generated, a loop in (Figure 5) will lead us back to selection, which is then followed by crossover, mutation and election before the next belief is generated. The loop will continue until all J beliefs of $\{B^i_{j,t}\}^J_{j=1}$ are generated. One of the beliefs, $B^{i,*}_{j,t}$, will be chosen based on the likelihood criteria,

$$B^{i,*}_t = \arg\max_j L(\{m_{t-s}\}^v_{s=1} \mid B^i_{j,t}), \tag{28}$$

The belief set will remain unchanged for the next Δ periods, when another loop of revision and renewal process is conducted, and $B^{i,*}_{t+\Delta}$ is brought about.

Except for the *CAPM believers*, the procedure will be applied to all investors so that they update their beliefs.

3.4 The Behavior of CAPM believers

The agent-based artificial stock market can be used to *directly* test the survivability of different portfolio rules. To achieve this goal, we introduce the *CAPM*

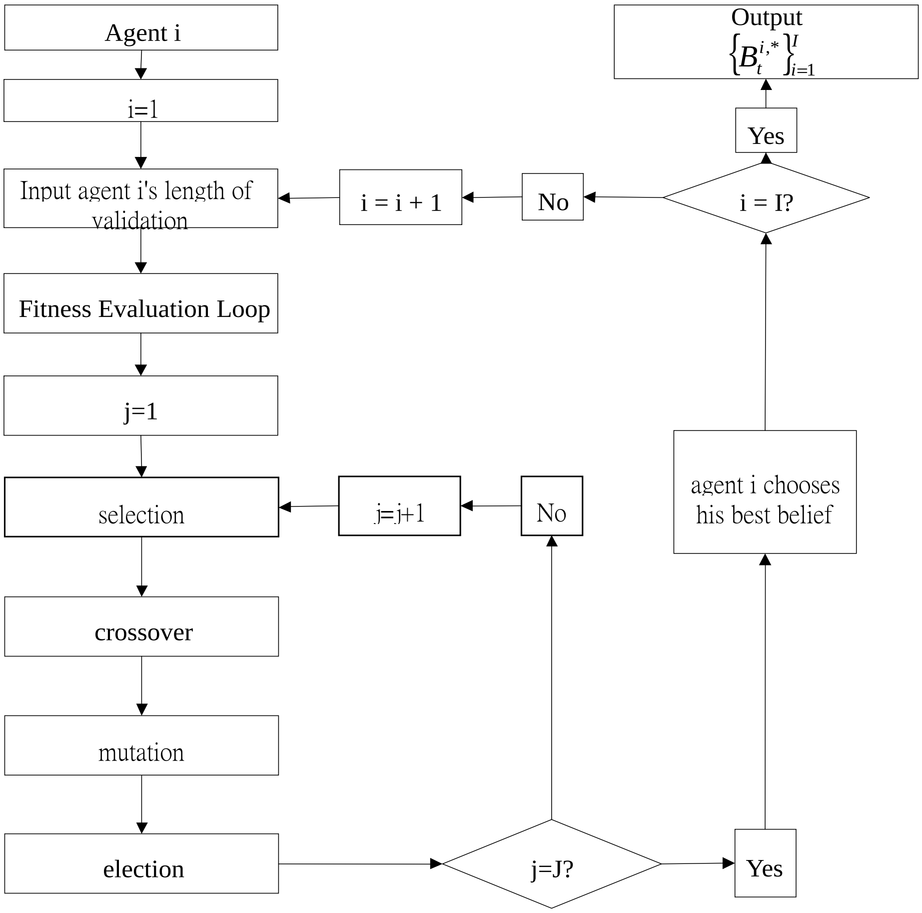

Fig. 5. Flowchart of the High-Level GA

believers, who use a prespecified formula (a portfolio rule) to determine their portfolios, and the portfolio behavior is time-invariant and not affected during the course of evolution.

The CAPM believers are investors who believe in the CAPM (type CAPM) and use it as a rule of thumb. They first work out the composition of the market and the risk-free portfolios. Then, according to their degree of risk aversion, they choose their preferred combination between the two. We index CAPM traders by means of κ; at date t, investor κ randomly chooses $\gamma_t^\kappa \in [0,1]$ from U(0,1) and invests in asset m a portion $\alpha_{m,t}^{CAPM(k)}$ of his savings such that:

$$\alpha_{m,t}^{CAPM(k)} = \gamma^k \alpha_{m,t}^F + (1-\gamma^k)\alpha_{m,t}^M \tag{29}$$

where $\alpha^F_{m,t} \equiv \frac{\hat{\rho}_{m,t}/w_{m,t}}{\sum_v \hat{\rho}_{v,t}/w_{v,t}}$ and $\alpha^M_{m,t} \equiv \frac{\hat{\rho}_{m,t}}{\sum_v \hat{\rho}_{v,t}}$. For simplicity, we assume that they have static expectations, i.e. $\hat{\rho}_{m,t} = \rho_{m,t-1}$ and $\hat{\rho}_{v,t} = \rho_{v,t-1}$.

As to their formula for savings rates, we design that in order to reflect the spirit of the formula for the CAPM portfolio as

$$\delta_t^{CAPM(k)} = \gamma^k \delta_t^F + (1-\gamma^k)\delta_t^M \tag{30}$$

where $\delta_t^F \equiv 0$ and $\delta^M_{m,t} \equiv \sum_{i=1}^{I} WS^i_{t-2} \cdot \delta^i_{t-1}$.

3.5 Summary

Figure 6 is a summary of the agent-based artificial stock market.

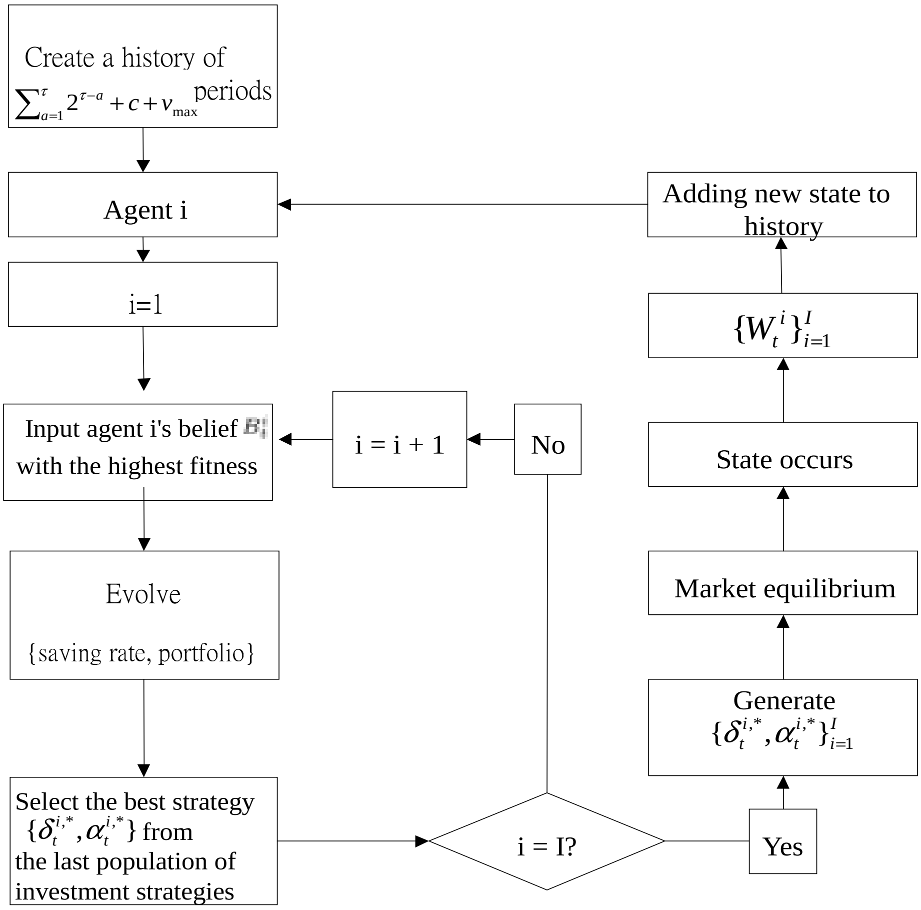

Fig. 6. A Summary of the Agent-Based Artificial Stock Markets

4 Simulation Results

4.1 Parameter Setup

The simulations we conducted share common values for the following parameters. There are 40 agents in the market, including 5 CAPM believers and 35 agents with 7 different types of utility functions. Therefore, $I = 40$. The 35 agents include 7 types of preference and one type includes 5 agents. The last type of utility function is generated from a utility generation program which generates 5 sets of parameters $(a_0, a_1, a_2, a_3, a_4, a_5, a_6)$ to make the 5 utility functions $u(c) = a_0 + a_1 \cdot c + a_2 \cdot c^2 + a_3 \cdot c^3 + a_4 \cdot c^4 + a_5 \cdot c^5 + a_6 \cdot c^6$ all satisfy the criteria that $u^{'} > 0$ and $u^{"} < 0$. The first 6 types of utility function are $\frac{1}{\beta_1} log(\alpha_1 + \beta_1 c)$, $\sqrt{c}$, $\alpha_3 + \beta_3 c$, $\frac{\alpha_4}{\beta_4} exp(\beta_4 c)$, $\frac{1}{(\gamma_5+1)\beta_5}(\alpha_5 + \beta_5 c)^{\gamma_5+1}$, $c - \frac{\alpha_6}{2} c^2$, respectively, where $\alpha_i, \beta_i, \gamma_i, i = 1, 3, 4...6$, are parameters, and c is the consumption amount.[2] These 6 types of agents will be called type 1, type 2...type 6 agents, respectively. The number of beliefs in each agent's belief set, J, is set to be 100. There are only 5 assets available in the market, and the dividends of assets 1, 2, 3, 4, 5 are 5, 4, 3, 2, 1, respectively. The number of investment rule that each agent can choose from, N, is 100. At each point of time, the investor has a perception of a time horizon with length $H = 25$. Agents simulate 5 H-horizon histories of the states based on their belief in order to evaluate the fitness of their investment rules; that is, $L = 5$. The number of generations that the *Low-Level GA* runs in one period, G, is set to be 50. The γ values exogenously given to 5 CAPM believers are 0.1, 0.2, 0.3, 0.4 and 0.5, respectively. The crossover rate is set to be 1 while the mutation rate is set to be 0.03.

We conducted 4 series of simulations as follows:

Table 1.

Parameters of utility function	Randomly generated		Exogenously given	
Agents share equal lengths of validation period	Series1:	100 times	Series2:	100 times
Agents set different lengths of validation period	Series3:	20 times	Series4:	20 times

In series 1 and 3, the parameters are randomly generated from 1 to 100. In series 2 and 4, we set the parameters : $(\alpha_1 = 0, \beta_1 = 1, \beta_3 > 0, \alpha_4 > 0, \beta_4 < 0, \alpha_5 > 0, \beta_5 > 0, \gamma_5 < 0, \alpha_6 < \frac{1}{5})$.

4.2 Main Results

Two main results are obtained. The first result is concerned with the beliefs. One phenomenon is that, through the calculation of the Kolmogorov-Smirnov

[2] Also see Chi-fu Huang and Robert H. Litzenberger (1988), pp.27-33

statistics, we find that agents' beliefs get closer to the true process in simulations assuming that agents have longer validation periods. Table 2, which follows, exhibits all agents' averages of the Kolmogorov-Smirnov statistics[3] representing each agent's accuracy of lifetime average belief under iid as the true model, and Table 3 exhibits the Markov process as the true model.

Table 2.

	v=200	v=100	v=50	v=30	v=15
1st simulation	0.08003069	0.061986	0.172932	0.304092	0.374112
2nd simulation	0.02486834	0.044006	0.165943	0.218958	0.297478
3rd simulation	0.06862377	0.09044	0.259228	0.311999	0.392694
4th simulation	0.05511849	0.06775	0.218943	0.294641	0.411535
5th simulation	0.03167815	0.018754	0.20938	0.340222	0.427986
6th simulation	0.05522913	0.056192	0.110966	0.174268	0.38312
7th simulation	0.07634258	0.059773	0.10703	0.0707	0.084455
8th simulation	0.07749052	0.047162	0.270188	0.357431	0.421027
9th simulation	0.05525984	0.089279	0.227164	0.300819	0.331442
10th simulation	0.04741502	0.058663	0.162812	0.227908	0.261784
Average	0.05720565	0.0594	0.190458	0.260104	0.338563

We may notice that, if the length of the validation is set to be longer, the accuracy of agents' beliefs will be enhanced.

Another phenomenon is that, regardless of whether the true model is iid or first-order Markov process, in different simulations with the same setting of parameters, each agent's belief regarding the dependence structure, k, may be 0 or 1, and it will switch over time.

The second result is concerned with the survivors. In our results for the series 1 simulations, except for the CAPM believers with low γ, the survivors are always type 2 agents. In our results for the series 2 simulations, by setting $\alpha_1 = 0, \beta_1 = 1$, except for the CAPM believers with low γ, the survivors are always type 1 and 2 agents. Therefore, we find that the utility function does play an important role in survivability. To enlarge the difference in terms of the accuracy of lifetime beliefs between agents in order to ensure that the belief is not as important as the preference, in series 3 and 4, we give the agents of the same type different lengths of validation, say 150, 100, 70, 50, and 30, and each type of agents share these levels of length. We find that the survivors are always type 2 agents with longer lengths of validation in series 3, and type 1 and 2 agents with longer lengths of validation in series 4. Hence, we conclude that the

[3] The Kolmogorov-Smirnov statistics calculate the largest difference in terms of the cdf between the belief and the true model.

Table 3.

	v=200	v=100	v=50	v=30	v=15
1st simulation	0.213899	0.208188	0.418389	0.506973	0.547073
2nd simulation	0.09156	0.099372	0.243963	0.313968	0.397163
3rd simulation	0.219324	0.210142	0.441415	0.558714	0.603704
4th simulation	0.036714	0.093825	0.193324	0.250146	0.325806
5th simulation	0.055175	0.084149	0.239975	0.348091	0.413207
6th simulation	0.147828	0.194238	0.377143	0.346744	0.399653
7th simulation	0.087423	0.132142	0.279255	0.316447	0.32759
8th simulation	0.118235	0.117335	0.225133	0.270151	0.335338
9th simulation	0.203246	0.196702	0.24599	0.243723	0.310246
10th simulation	0.05436	0.075059	0.179239	0.241425	0.282946
Average	0.122776	0.141115	0.284383	0.339638	0.394272

preference plays a more important role in survivability while the accuracy of beliefs still matters.

5 Discussions

5.1 Discussions about the beliefs

The first phenomenon is easy to understand for, when the validation period is longer, the state pattern of the validation period will be closer to the true model arising from the larger sample. This leads agents to choose to observe a longer period for the training data, i.e. a larger d, because the training data are also generated from the true model.

As to the second phenomenon, in this belief scheme, agents will try to find the (k, d) whose probability function summarized from the training data can fit the state pattern of validation period best. However, the data in the validation period are finite, and their pattern is impossible to perfectly match the true process. In further discussions which follow below, we shall discuss the reason why agents sometimes choose to believe that $k = 1$, i.e. that the Markov process under iid is the true model, and the reason why agents sometimes choose to believe that $k = 0$ under the assumption that the Markov model is the true model.

Under iid being the true model We first discuss the reason why agents may choose to believe that $k = 1$ under iid as the true model. Let us take one simulation with v=200 as an example. At time t=100, a typical belief with a markov transition matrix is shown as Table 4.

Furthermore, the true iid process is (0.21037, 0.20731, 0.10366, 0.18903, 0.28963). As we can see, each row of this belief is similar to the true model.

Table 4.

t \ t+1	1	2	3	4	5	K-S statistics
1	0.25	0.19318	0.07955	0.18182	0.29545	0.03963
2	0.2	0.25	0.13	0.23	0.19	0.09963
3	0.16071	0.19643	0.125	0.21429	0.30357	0.06054
4	0.19101	0.24719	0.13483	0.19101	0.23596	0.05367
5	0.15447	0.20325	0.13821	0.17074	0.33333	0.05996

It is impossible for the state pattern in the validation period to be exactly the same as the true one unless the length of the validation period is infinite. Therefore, agents may use the Markov table to approximate the true table. In addition, by deviating the probabilities of each row vector from the true iid probability vector, the Markov table may fit the pattern of the validation period even better.

Under Markov being the true model If the true model follows a Markov process, agents may choose to believe that state occurrence follows the iid process. This is not surprising because they are trying to fit the stationary distribution of the true Markov model. To examine this, let us take some simulation with its length of validation being 100 as an example. The stationary distribution of the true Markov table is (0.12513, 0.31706, 0.08811, 0.19841, 0.27233). Furthermore, the number of occurrence times for each state in the 100 periods is (7, 38, 9, 20, 26), respectively. As we can see above, the frequency of occurrence of each state is similar to the stationary distribution of the true Markov model. Now let us take a belief with k=0 as an example. Its corresponding probability function summarized from the training data at that time point is (0.11944, 0.28238, 0.07742, 0.24413, 0.27812). In addition, it is similar to the stationary distribution of the true Markov model. Therefore, we can infer that the reason why this agent tries this iid vector is that he is trying to fit the stationary distribution of the true Markov model.

Therefore, the survival pressure derived from the finite sample does not necessarily guarantee the convergence of the belief to the true model. However, this answer is not atypical. In fact, it is similar to what [Lettau (1997)] and many others found.

5.2 Discussions about the survivors

As we described, in series 1 simulations, type 2 agents and CAPM believers with low γ exhibit the strongest survivability; in series 2 simulations, type 1 and 2 agents and CAPM believers with low γ exhibit the strongest survivability. Therefore, we shall now discuss the survival factors of type 1 and 2 agents and CAPM believers with low γ, respectively.

The survival factor of type 2 agents The reason for their strong survivability is nothing to do with their beliefs. This is because it is impossible for certain agents' average lifetime beliefs to always be the most accurate in several different simulations. In other words, the reason why agents with the squared-root (c) utility always survive is not likely to be that their lifetime beliefs always happen to be the most accurate. In fact, we take some typical simulation of series 1 in which type 2 agents are the only survivors except for CAPM believers as an example. We compute the average cdf for each agent's lifetime beliefs, and compute the largest distance when compared with the true distribution, the Kolmogorov-Smirnov statistics. Then we average the Kolmogorov-Smirnov statistics of 5 agents with the same utility, to obtain the Kolmogorov-Smirnov statistics of 7 utility types which are: (0.01866, 0.01925, 0.01905, 0.01867, 0.01941, 0.01818, 0.01888). Obviously, the Kolmogorov-Smirnov statistics for type 2 agents, 0.01925, is not the smallest one. That is, their lifetime belief is not the most accurate one. However, they survive in this simulation.

Therefore, the only possible reason is that the investment rule derived from maximizing the expected discounted square-root (c) utility function carries some characteristics that support their survival. The investment rule includes two parts, one being the portfolio rule and another is the saving rate. From several examples we find that the portfolio rule of agents with the squared-root (c) utility is not better then that of others. Taking the same case above as an example, we first compute each agent's lifetime average portfolio, $(AP_1^i, AP_2^i, AP_3^i, AP_4^i, AP_5^i)$. Then we compute the expected payoff: $EP^i = AP_1^i * Ew_1 + AP_2^i * Ew_2 + \ldots + AP_5^i * Ew_5$ for each agent i's lifetime average portfolio, where $Ew_m, \forall m = 1...5$ is the probability of a state m occurrence times the dividend of asset m. Then we average the EP value of 5 agents with the same utility, and we obtain the averaged EP for 7 utility types: (0.611386, 0.609724, 0.615229, 0.609119, 0.647657, 0.645967, 0.603597). As we can see, the average EP of agents with the squared-root (c) utility, 0.609724, is not the highest one. Therefore, the portfolio rule for agents with the squared-root (c) utility is not the best in this case. However, they indeed survive.

This forces us to notice the characteristics of saving rates that agents choose. Furthermore, an obvious characteristic has been observed to support the argument that the saving rate plays a much more important role than the portfolio in determining the survival of agents. The obvious characteristic is that the saving rate chosen by agents with square-root (c) utility is most stable for life. Unlike other agents, they never choose a saving rate that is too low in their life. We believe that this is the main reason for their survival.

Let us take 10 typical simulations of series 1 (5 simulations are conducted under the true model being iid, and 5 under the markov process. All of them are simulated under a length of validation period of 100) in which agents with square-root(c) utility survive with some CAPM believers as an example. For each example, we average the maximum, higher quantile point, medium, lower quantile point and minimum of their saving rates over the 100 periods for 5 agents with the same utilty. Then, we average the values in 10 simulations, and obtain Table 5 in which each column shows the values of each utility type, respectively.

Table 5.

min	0.0027	0.20773	0.0019	0.0011	0.0026	0.05245	0.0016
Q1	0.2935	0.42391	0.23788	0.2387	0.4979	0.11594	0.49093
medium	0.575	0.61456	0.91318	0.44552	0.92016	0.72471	0.57388
Q3	0.72688	0.90163	0.95569	0.84291	0.95553	0.94941	0.88032
max	0.96668	0.96034	0.95968	0.96156	0.9685	0.96157	0.96918
mean	0.246825	0.391177	0.406327	0.248742	0.439013	0.297095	0.286658

As we can see, the mean of the saving rates of agents with the squared-root (c) utility, 0.39118, is not the highest. However, they never choose an extremely low saving rate. Figure 7 is the corresponding Box-Whisker plot. The horizontal axis represents the different utility types, and the vertical axis represents the levels of saving rates.

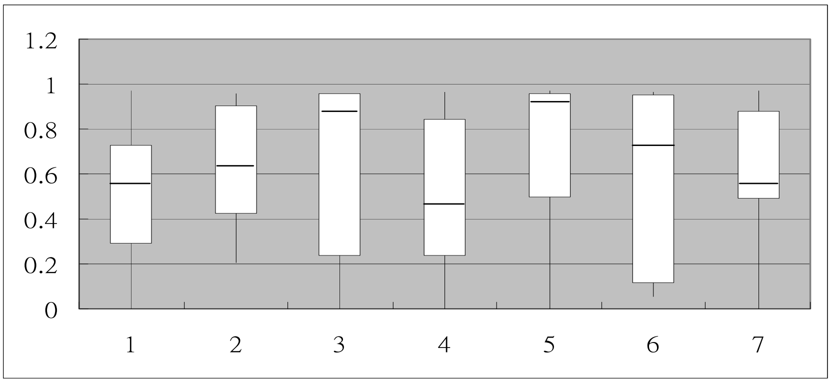

Fig. 7. Box-Whisker plot of saving rates

Now let us observe and discuss these points in more detail. We notice that every time when the asset prices for the last period are particularly high on average, the agents' saving rates decrease for this period. This is because they use the prices of the last period as expected prices to estimate the return from the investing their J rules as described in Section 3.1; therefore, when the prices of the last period are on average high, they predict that the investing will bring a low return and, therefore, they will tend to choose the investment rule with a low saving rate.

Although all agents lower their saving rates, type 2 agents react mildly compared to others. Take a typical simulation as an example. In this simulation, the average prices of periods 10, 12, 18 and 20 are particularly high, and the saving rates of agents, excluding type 2 agents, for periods 11, 13, 19 and 21 decrease dramatically. Table 6 shows the saving rates for different types of agents in the first 25 periods.

Table 6.

	type 1	type 2	type 3	type 4	type 5	type 6	type 7
period 1	0.794	0.919	0.959	0.8405	0.953	0.4023	0.9566
period 2	0.267	0.3036	0.006	0.006	0.55	0.32	0.049
period 3	0.828	0.9522	0.954	0.869	0.961	0.3225	0.952
period 4	0.8519	0.679	0.917	0.944	0.961	0.8256	0.9578
period 5	0.8519	0.679	0.9327	0.944	0.96	0.8226	0.9578
period 6	0.017	0.463	0.0347	0.015	0.96	0.2964	0.067
period 7	0.8957	0.855	0.958	0.9457	0.96	0.7101	0.959
period 8	0.051	0.396	0.0142	0.1509	0.6511	0.024	0.013
period 9	0.937	0.8614	0.9626	0.949	0.9638	0.956	0.9503
period 10	0.8097	0.5839	0.7807	0.5977	0.9638	0.6641	0.889
period 11	0.005	0.4423	0.006	0.001	0.004	0.003	0.0729
period 12	0.917	0.6024	0.8795	0.8768	0.9562	0.872	0.6662
period 13	0.001	0.2109	0.001	0.003	0.007	0.001	0.001
period 14	0.8869	0.9556	0.956	0.8906	0.951	0.052	0.9567
period 15	0.1179	0.3861	0.007	0.01	0.004	0.153	0.0355
period 16	0.9532	0.9104	0.951	0.949	0.9482	0.956	0.951
period 17	0.9566	0.926	0.957	0.9561	0.9605	0.956	0.959
period 18	0.9566	0.926	0.9552	0.9561	0.9605	0.956	0.958
period 19	0.006	0.241	0.003	0.001	0.007	0.002	0.006
period 20	0.943	0.9422	0.9653	0.9418	0.9591	0.957	0.956
period 21	0.009	0.2209	0.073	0.008	0.1601	0.059	0.036
period 22	0.963	0.9535	0.959	0.9601	0.954	0.96	0.9375
period 23	0.9184	0.6517	0.9345	0.905	0.9235	0.9444	0.907
period 24	0.037	0.3027	0.07	0.453	0.002	0.001	0.619
period 25	0.9564	0.9486	0.95	0.9643	0.9592	0.959	0.9548

We also observe that, when the period that an agent chooses an extremely low saving rate, his wealth decreases terribly. Furthermore, after this kind of event happens several times, he is driven out of the market. Therefore, we conclude that the stablility of saving rates is the most important factor in terms of the survival of agents with square-root (c) utility functions.

The survival factor of type 1 agents In series 2 simulations in which we set $\alpha_1 = 0, \beta_1 = 1$, we found that type 1 agents show strong survivability, too. We observed their saving rates, and found that, over the 100 periods, the saving rates became nearly fixed at their discount rate, the β value, along the time path.

[Blume and Easley (1992)] provide the analytical solution for the investment rule maximizing expected discounted utility in Theorem 5.1 that states “Suppose trader i's objective function is $E^i[\sum_{t=1}^{\infty} log(c_t^i)]$. If beliefs over states at date t are q_t^i and the value above is finite for any investment rule, then the optimal investment rule is the simple rule $\delta_t^i = \beta^i, \alpha_t^i = q_t^i$ at each date."

Our *Low-Level GA* does the optimality job quite well. Take some type 1 agent in a simulation as an example. His lifetime average investment rule ($\bar{\delta^i}, \bar{\alpha_1^i}, \bar{\alpha_2^i}, \bar{\alpha_3^i}, \bar{\alpha_4^i}, \bar{\alpha_5^i}$)=(0.5897, 0.0589, 0.1847, 0.3006, 0.2484, 0.2074) approximates the theoretically optimal rule ($\beta, \bar{q_1^i}, \bar{q_2^i}, \bar{q_3^i}, \bar{q_4^i}, \bar{q_5^i}$)=(0.59, 0.0476, 0.1898, 0.3002, 0.2502, 0.2107).

Similarly, we could find some examples in which they survive while their portfolio rule and the accuracy of their lifetime beliefs are not the best among agents. Again, the stablility of the saving rates is the key factor in their survival.

The reason for type 1 and type 2 agents choosing stable saving rates From the *Euler equation*, the basic condition for choosing consumption over time is known: $r = \beta - [\frac{u"(c) \cdot c}{u'(c)}](\frac{\dot{c}}{c})$.

where r is the rate of return on investment.

Because the coefficient of RRA (Relative Risk Aversion) is defined as $-[\frac{u"(c) \cdot c}{u'(c)}]$, the *Euler equation* described above can be rewritten as follows:

$$r = \beta + RRA \cdot (\frac{\dot{c}}{c}) \tag{31}$$

From Equation (31), we know that when r decreases, if someone's coefficient of RRA approaches zero, his $\frac{\dot{c}}{c}$ must drop dramatically. This will drive agents with the coefficients of RRA approaching zero to choose an extremely low saving rate when the rate of return on investment clearly decreases.

Table 7 summarizes the coefficient of RRA for each type of agent in series 1 and 2.

Obviously, in series 1, the only agents with constant RRA are type 2 agents (RRA=0.5). This CRRA characteristics free them from the crisis of zero RRA approaching and then win them a stable saving rate path. The reason for the stable saving rate path of type 1 agents in series 4 is similar. Their utility functions are CRRA (RRA=1), too.

The survival factor of CAPM believers with low γ values CAPM believers do not have to guess the true distribution, but instead follow the formula described earlier.

Notice that the market portfolio part (and the corresponding part in the savings rule) reflects the weighted average of other agents' investment rules (in this period if they have perfect foresight). In particular, it assigns a larger weight to the dominant agents' portfolio. If his γ value is 0, then the CAPM believer behaves just as a dominant agents' imitator. Hence, the market portfolio part is the superior position of their rules.

We further analyze the reason why being endowed with a high γ value is not good for survival as follows. In the risk-free portfolio part, the original idea

Table 7.

utility type	RRA	in series 1	in series 2
$u(c) = \frac{1}{\beta_1} \log(\alpha_1 + \beta_1 c)$	$\frac{1}{\frac{\alpha_1}{\beta_1 c} + 1}$	approaches 0	1
$u(c) = \sqrt{c}$	0.5	0.5	0.5
$u(c) = \alpha_3 + \beta_3 c$	0	0	0
$u(c) = \frac{\alpha_4}{\beta_4} \exp\{\beta_4 c\}$	$-\beta_4 c$	approaches 0	approaches 0
$u(c) = \frac{1}{(\gamma_5 + 1)\beta_5} (\alpha_5 + \beta_5 c)^{\gamma_5 + 1}$	$-\frac{\beta_5 \gamma_5}{\frac{\alpha_5}{c} + \beta_5}$	approaches 0	approaches 0
$u(c) = c - \frac{\alpha_6}{2} c^2$	$\frac{\alpha_6}{\frac{1}{c} - \alpha_6}$	approaches 0	approaches 0
$u(c) = a_0 + a_1 c + a_2 c^2 + a_3 c^3 + a_4 c^4 + a_5 c^5 + a_6 c^6$	$-\frac{2a_2 c + 6a_3 c^2 + 12a_4 c^3 + 20a_5 c^4 + 30a_6 c^5}{a_1 + 2a_2 c + 3a_3 c^2 + 4a_4 c^3 + 5a_5 c^4 + 6a_6 c^5}$	approaches 0	approaches 0

The reason why some types of agents' RRA approaches zero is that the wealth on average is quite small, not to mention the consumption. When c approaches zero, those RRA values also approach zero.

of dividing the asset price by its dividend is that the risk of the asset with a higher dividend is always higher, so investing assets with high dividends are not suggested from the risk-free point of view. However, in our economy, the worst return from investing in any asset is just zero. This risk-free portfolio part makes CAPM believers invest less in assets that give higher dividends than the optimal portfolio. Hence, we think that the risk-free portfolio part is an inferior factor for CAPM believers to survive in our economy.

Besides, agents in the economy are assumed to be boundedly-rational. We further assume that they have static expectations, i.e. $\hat{\rho_{m,t}} = \rho_{m,t-1}$. Therefore, what CAPM believers mimic is other agents' portfolios in the previous period. However, this only misleads CAPM believers in one kind of situation. Let us explain this step by step.

There are two factors that determine non-formula agents' portfolios. One is the assets' dividends, but they are time-invariant. Another is the belief regarding the probability that each asset will give a dividend, and the belief is time-variant.

When the true model follows the iid process, the probability that each state will occur will never change from the beginning to the end. If agents believe that the true model follows the iid process, they will try to fit **the same** model in each period. Therefore, it does not make much difference if the CAPM believers mimic non-formula agents' beliefs in the last or current period. Even if they believe that true model follows a Markov process, each row in their Markov table approximates the same time-invariant true iid probability vector. Therefore, the CAPM believers' mimicking the non-formula traders' beliefs in different periods is the same as mimicking similar probability functions in different rows. Therefore, when the iid is the true model, the effect of boundedly-rational CAPM believers mimicking other agents' investment rules in the last period is similar to the effect where perfect foresight is assumed.

When the true model follows the Markov process, if other agents' beliefs are trying to fit the stationary distribution of the true model, the effect of boundedly-rational CAPM believers mimicking other agents' last-period portfolios is the same as the effect where perfect foresight is assumed, because the stationary distribution will also never change from beginning to the end. However, if agents believe that the Markov process is the true model, boundedly-rational CAPM believers will be misled. Take two of our simulations for a comparison. We assume that the true model follows the Markov process, that the lenghth of the validation period is 100, and that non-formula agents' saving rates are equal to 0.59 in both cases. The γ values assigned to each set of CAPM believers are 0.1, 0.2, 0.3, 0.4, 0.5 in both cases. The only difference is that, in the first case, most of the time agents believe that the true model follows the iid process and, in the second case, the markov process. The CAPM believer with the γ value =0.1 survives in both cases. However, the agent with γ value=0.2 can only survive in the first case. Therefore, we believe that CAPM believers are more or less misled in the second situation, and hence, only the CAPM believers with the lowest γ value can survive. This finding tends to support our hypothesis.

6 Conclusion

Our agent-based simulation results are largely consistent with [Blume and Easley (1992)]. First, we also find that rational log-utility traders survive in series 2 simulations.[4] Secondly, rational CRRA agents with moderately high RRA coefficients such as type 2 agents also survive. In fact, we also try $u(c) = c^{\alpha_2}$, and we explore several levels of parameter α_2 values. We find that when α_2 increases (meaning that the agent's coefficient of RRA decreases), his survivability decreases. Thirdly, forecasting accuracy does not guarantee survival. The example mentioned in Section 5.2 shows that, despite their lowest lifetime Kolmogorov-Smirnov statistics, 0.01818, the type 6 agents are driven out in that simulation. Furthermore, to enlarge the differences in the accuracy of lifetime beliefs among agents, in series 3 and 4, we set that the same types of agents have different lengths of validation, say 150, 100, 70, 50 and 30, and that each type of agents shares these levels of length. We find that the survivors are always type 2 agents with lengths of validation equal to 150 and 100 (70 also survive sometimes) in series 3, and type 1 and 2 agents with length of validation equal to 150 and 100 (70 also survive sometimes) in series 4. The survivability of other types of agents with a length of validation equal to 150 is worse than that of type 2 agents with a length of validation equal to 70 in series 3. Hence, we conclude that preference plays a more important role in survivability, although accuracy of belief does still matter.

[4] The type 1 agents in series 1 have a different coefficient of RRA from that of the log utility function discussed in [Blume and Easley (1992)].

However, in using the agent-based model, we also find something which was not shown in [Blume and Easley (1992)]. First, there are other types of traders that may survive in the market as well, e.g., the CAPM traders. Secondly, the reason why CRRA traders can survive in the market is because of their implied stable saving behavior. This may help us understand locked-up saving contracts, such as the national annuity program. Finally, if we assume that agents have the same discount factor, the result of [Blume and Easley (1992)] remains robust even if saving rates are determined endogeously. Furthermore, their savings rule also exhibits superiority in determining survival.

References

[Arthur et. al. (1997)] Arthur, W.B., Holland, J., LeBaron, B., Palmer, R. and Tayler, P. (1997), "Asset Pricing under Endogenous Expectations in an Artificial Stock Market," In: Arthur, W.B., Durlauf, S. and Lane, D. (Eds.), *The Economy as an Evolving Complex System II*, Addison-Wesley, Reading, MA, pp. 15–44.

[Blume and Easley (1992)] Blume, E. and E. Easley (1992), "Evolution and Market Behavior," *Journal of Economic Theory*, Vol. 58, pp. 9-40.

[Cacho and Simmons (1999)] Cacho O., and P. Simmons (1999), "A Genetic Algorithm Approach to Farm Investment," *Australian Journal of Agricultural and Resource Economics*, Vol. 43, No. 3, pp. 305–322.

[Chen and Yeh (2001)] Chen S.-H., and C.-H. Yeh (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*, Vol. 25, pp. 363–393.

[Chen and Yeh (2002)] Chen S.-H., and C.-H. Yeh (2002), "On the Emergent Properties of Artificial Stock Markets: The Efficient Markets Hypothesis and the Rational Expectations Hypothesis," *Journal of Economic Behavior and Organization*, forthcoming.

[De Long et al. (1991)] De Long, J., A. Shleifer, L. Summers, and R. Waldman (1991), "The Survival of Noise Traders," *Journal of Business*, Vol. 64, pp. 1-19.

[Farmer and Lo (1999)] Farmer, J.D. and Lo, A.W. (1999), "Frontiers of Finance: Evolution and Efficient Markets," *Proceedings of the National Academy of Sciences*, Vol. 96, pp. 9991–9992.

[Kelly (1956)] Kelly, J. L. (1956), "A New Interpretation of Information Rate," *Bell System Technical Journal*, Vol. 35, pp. 917-926.

[Holland and Miller (1991)] Holland, J., and J. Miller (1991), "Artificial Adaptive Agents in Economic Theory," *American Economic Review*, Vol. 81, No. 2, pp. 365–370.

[Jagannathan and Wang (1996)] Jagannathan, R. and Z. Wang (1996), "The Conditional CAPM and the Cross-Sections of Expected Returns," *Journal of Finance*, Vol. 51, pp. 3-53.

[LeBaron et al. (1999)] LeBaron B., W. B. Arthur, and R. Palmer (1999), "Time Series Properties of an Artificial Stock Market," *Journal of Economic Dynamics and Control*, Vol. 23, pp. 1487–1516.

[LeBaron (2000)] LeBaron, B. (2000), "Agent Based Computational Finance: Suggested Reading and Early Research," *Journal of Economic Dynamics and Control* Vol. 24, pp. 679–702.

[LeBaron (2001)] LeBaron, B. (2001), "Evolution and Time Horizons in an Agent Based Stock Market," *Macroeconomic Dynamics*, forthcoming.

[Lettau (1997)] Lettau, M. (1997), "Explaining the Facts with Adaptive Agents: the Case of Mutual Fund Flows," *Journal of Economic Dynamics and Control*, Vol. 21, No. 7, pp. 1117–1147.

[Merton and Samuelson (1974)] Merton, R. C. and P. A. Samuelson (1974), "Fallacy of the Log-Normal Approximation to Optimal Portfolio Decision-Making over Many Periods," *Journal of Financial Economics*, Vol. 1, pp. 67-94.

[Sandroni (2000)] Sandroni, A. (2000), "Do Markets Favor Agents Able to Make Accurate Predictions?" *Econometrica*, Vol. 68, No. 6, pp. 1303-1341.

[Sciubba (1999)] Sciubba, E. (1999), "The Evolution of Portfolio Rules and the Capital Asset Pricing Model", DAE Working Paper no. 9909, University of Cambridge.

[Spear (1989)] Spear, S. (1989). "Learning Rational Expectations Under Computability Constraints," *Econometrica* Vol. 57, pp. 889-910.

[Szpiro (1997)] Szpiro, G. G., (1997), "The Emergence of Risk Aversion," *Complexity* 2, pp. 31–39.

[Tay and Linn (2001)] Tay, N. and S. Linn (2001), "Fuzzy Inductive Reasoning, Expectation Formation and the Behavior of Security Prices," *Journal of Economic Dynamics and Control*, Vol. 25, pp. 321-361.

[Tesfatsion (2001)] Tesfatsion, L. (2001), "Introduction to the Special Issue on Agent-Based Computational Economics," *Journal of Economic Dynamics and Control*, Vol. 25, pp. 281–293.

[Vriend (2000)] Vriend, N. J. (2000), "An Illustration of the Essential Difference between Individual and Social Learning, and its Consequences for Computational Analyses," *Journal of Economic Dynamics and Control*, Vol. 24, pp. 1–19.

[Litzenberger (1988)] Huang, C. F. and Litzenberger, R. H. (1988), *Foundations for Financial Economics*, Prentice Hall, pp. 27–33.