

Risk Preference, Forecasting Accuracy and Survival Dynamics: Simulations Based on a Multi-Asset Agent-Based Artificial Stock Market

Shu-Heng Chen ^{a,*}, Ya-Chi Huang ^b

^a*AI-ECON Research Center, Department of Economics, National Chengchi University, Taipei, Taiwan 116*

^b*AI-ECON Research Center, Department of Economics, National Chengchi University, Taipei, Taiwan 116*

Abstract

The relevance of risk preference and forecasting accuracy to the survival of investors is an issue that has recently attracted a number of recent theoretical studies. At one extreme, it has been shown that risk preference can be entirely irrelevant, and that in the long run what distinguishes the agents who survive from those who vanish is just their forecasting accuracy. Being in line with the market selection hypothesis, this theoretical result is, however, established mainly on the basis of Pareto optimal allocation. By using agent-based computational modeling, this paper extends the existing studies to an economy where adaptive behaviors are autonomous and complex heterogeneous, and where the economy is notorious for its likely persistent deviation from Pareto optimality. Specifically, a computational multi-asset artificial stock market corresponding to Blume and Easley (1992) and Sandroni (2000) is constructed and studied. Through simulation, we present results that contradict the market selection hypothesis. Among the eight types of agents considered in this model, only log-utility agents survive, and the rest are driven out, including even those who have superior forecasting accuracy. Nevertheless, when all the agents are of the same type, the wealth share is positively correlated to forecasting accuracy, and the market selection hypothesis is sustained, at least in a weak sense.

Key words: Market selection hypothesis, Agent-based artificial stock markets, Autonomous agents, Genetic algorithms

* Corresponding author. Tel/fax: +886-2-29387308.

Email addresses: chchen@nccu.edu.tw (Shu-Heng Chen), yachi@aiecon.org (Ya-Chi Huang).

URL: <http://www.aiecon.org> (Shu-Heng Chen).

1 Introduction

Agent-based computational economic (hereafter **ACE**) modeling is distinguished from the conventional economic modeling by its great flexibility in terms of agents' *heterogeneity* and the associated *population dynamics*. This advantage may be very helpful in studying the *survivability* of different types of agents, specifically when they are placed in a complex interactive environment. In this paper, the ACE approach is applied to address a debate which can be related to the *market selection hypothesis*, according to which markets favor rational traders over irrational traders.

The debate, if we trace its origin, started as a result of the establishment of what become known as the *Kelly criterion* (Kelly, 1956), which basically says that a rational long run investor *should* maximize the expected growth rate of his wealth share, and therefore should behave as if he were endowed with a logarithmic utility function. In other words, the Kelly criterion implicitly suggests that there is an optimal preference (rational preference) which a competitive market will select and that is *logarithmic utility*. The debate on the Kelly criterion has a long history, and so not surprisingly, there is a long list of both pros and cons with regard to it as the literature develops.¹

A possible implication of the Kelly criterion is that an agent who maximizes his expected utility under the *correct* belief may be driven out by an agent who maximizes his expected utility under an *incorrect* belief, simply because the former does not maximize a logarithmic utility function, whereas the latter does. Blume and Easley (1992) were the first to show this implication of the Kelly criterion in a competitive asset market. In their seminal study, they questioned the survivability of rational investors. In a nutshell, they showed that *rational investors* who are characterized by their selection of a portfolio that maximizes their expected utility with respect to the correct belief may not be good enough to survive. To enhance their survivability, their preference over risk (utility function) must also be "optimal". If not, an even more striking result is that these rational agents may be driven out of the market by those agents who base their decisions on *incorrect beliefs*, but have a "nearly optimal" preference.²

The market selection hypothesis, therefore, fails because agents with accurate beliefs are not necessarily selected. A consequence of this failure is that asset prices may not eventually reflect the true value of the asset, and may fail to converge to the *rational expectations equilibrium*.

Nonetheless, a series of recent studies indicates that the early analysis of Blume and Easley (1992) is not complete. Sandroni (2000) shows that, if the saving behavior is endogenously determined, then the market selection hypothesis is rescued,

¹ See Sciubba (1999) for a quite extensive review.

² Other similar findings can also be found in Sciubba (1999).

and in the long-run, only those optimizing investors with *correct beliefs* survive. The surviving agents do not have to be log-utility maximizers, and they can have diverse preferences over risk. Sandroni's analysis is further confirmed by Blume and Easley (2001) in a connection of the market selection hypothesis to the *first theorem of welfare economics*. They show that in a dynamic complete market *Pareto optimality* is the key to understanding selection for or against traders with correct beliefs: in any optimal allocation the survival or disappearance of a trader is determined entirely by beliefs, and not by risk preference.

Sandroni (2000)'s and Blume and Easley (2001)'s studies are largely analytical. They both take a Pareto optimal allocation as a starting point to work with. The dynamic process converging to a Pareto optimal allocation itself is, nonetheless, beyond the scope of their analysis. Issues related to the dynamic process are two-fold. First, there is *individual dynamic optimization*. A Pareto optimal allocation rests upon the optimization of all individuals. In this specific context, this requires that all agents are able to solve the infinite-time stochastic dynamic optimization problem facing them, regardless of their preferences over risk or utility functions. However, analytical solutions known to us are severely restricted to certain classes of preferences. In general, one has to rely on numerical approximation, which means that Pareto optimality may not always be attainable.

What makes this problem even more complex is, however, the second issue: *trading at an equilibrium consistent with price expectations*. Notice that what we study here is not a simple representative-agent optimization problem, but a market composed of heterogeneous agents. Each one of them, upon maximizing his expected utility, has to know the prices of assets in the future. These prices are, nonetheless, endogenously generated by agents' own perceptions. As a result, a typical fixed-point problem occurs. The market, as a distributed decentralized processor, may fail to coordinate its participants to such a fixed point.³ In general, it will depend on agents' forecasting rules and the associated learning schemes, and it is likely that agents will trade at prices that are inconsistent with their ex-ante expectations of the prices. In this case, Pareto optimality is also not attainable.

Both of the two issues discussed above are directly related to the attainability of Pareto optimality. However, Pareto optimality per se was only taken by Sandroni and Blume and Easley as a convenient starting point for their analytical work. To facilitate their further analysis, the learning dynamics concerned with the updating of agents' beliefs also be needed to be simplified. Sandroni, for example, did not deal with learning dynamics directly; instead, he assumed that there will be a day when *some agents can eventually make accurate predictions* or *eventually make*

³ In contrast to the agent-based bottom-up approach, there is a top-down way to avoid this issue by using the design of a Walrasian auctioneer and simply equating demand to supply. However, deriving the aggregate demand function for each asset under different sets of prices itself is a daunting task.

accurate next period predictions, and started his major analytical work from there. Nevertheless, a plausible process to show the appearance of these *sages* was absent. It is, therefore, not entirely clear whether these types of agents will ever emerge.⁴ What happens when no trader has correct beliefs?⁵

Blume and Easley (2001) do recognize that the market selection hypothesis would be of little interest if it were to address only selection for traders with correct beliefs. Their delicate analysis of learning leads to two major findings as to the superiority of Bayesians. First, a Bayesian almost surely survives for almost all possible truths in the support of her prior. Second, in the presence of a Bayesian trader, any traders who survive are not too different from Bayesians. We admire the beauty of the analysis of the Bayesians, but are not entirely easy about traders being just Bayesians. For us, being a Bayesian is only one way of representing human learning behavior. It has its mathematical rigor, but when cognitive and computational constraints are presented, it is not clear whether Bayesian learning is the most effective way of describing observed learning behavior. The experimental evidences certainly not always in favor of Bayesian learning.⁶ Therefore, this consideration does not stop us from asking: what happens when no traders are Bayesians?

This review and discussion of the early literature now seems to indicate clearly where we are moving. Needless to say, the above-mentioned analytical work on the market selection hypothesis has already provided us with an interesting benchmark to reflect upon, namely, *the irrelevance of risk preference*. However, since the conclusive statement is very interesting, it would be useful to see how strongly we can put it by relaxing some tight constraints. In this paper, we do not assume Pareto optimality, the emergence of the sages, or the Bayesians. This relaxation allows for a more extensive class of bounded-rational behaviors, and we examine whether the irrelevance of risk preference still holds with this enlargement. Furthermore, given the popularity of experimental economics, this enlargement will also enable us to clothe the debate with some empirical or experimental content, on which a laboratory design or an empirical study can actually be based at a later stage.

Concretely speaking, what is proposed here is a computational model, namely, an agent-based computational version of Blume-Easley-Sandroni's model. This

⁴ These sages are rigorously described by Sandroni using concepts from probability theory, such as *merging* or *weakly merging*. However, these conditions are not easy to verify empirically when we would like to know whether the history of mankind has ever experienced such sages.

⁵ Sandroni (2000) does consider the case when no one has correct beliefs. His proposition 3 basically compares two kinds of agents: one persistently forecast more accurate than the other, while both do not have correct beliefs. He then shows that the former will drive out the latter. In Section 5.2, we shall design a specific experiment to test some related aspects of this proposition.

⁶ Feldman (1962) is one of the famous early examples, and many more can be found in the survey articles by Rabin (1998) and Barberis and Thaler (2002).

ACE platform enables us to examine the *survivability* of a type of agents when they are interacting with other types of agents in a real-time competitive environment. If the type of the agents happens to be an “incarnation” of an economic hypothesis, then the use of the ACE model can become a *survival test*, as an alternative to the econometric test and the laboratory test, for the hypothesis. Over the last few years, this survival test has been applied to different levels of economic behavior, such as different trading strategies, different forecasting models, different decision variables, but all under the given objectives. Here, we shall extend the survival test to a higher level, i.e. to examine the significance of the risk preference in determining survivability.

The rest of the chapter is organized as follows. Section 2 briefly reviews the Blume-Easley-Sandroni model. An agent-based computational version of the model is provided in Section 3, which can be regarded as an extension of the single-asset artificial stock market to its multi-asset version. The debate then proceeds with the experimental designs given in Section 4. The simulation results and analysis are provided in Section 5, followed by the concluding remarks in Section 6.

2 The Blume-Easley-Sandroni Model

Our agent-based artificial stock market is built upon the analytical model which was first initiated by Blume and Easley (1992) and was later extended by Sandroni (2000). The Blume-Easley-Sandroni (hereafter, the **BES**) model is briefly reviewed in this section.

Consider a complete securities market. Time is discrete and indexed by $t = 0, 1, 2, \dots$. There are M states of the world indexed by $m = 1, 2, \dots, M$, one of which will occur at each date. States follow a stochastic process. Asset m pays dividends $w_m > 0$ when state m occurs, and 0 otherwise. At each date t , the outstanding volume of each asset is exogenously fixed at one unit, so that the total wealth in the economy at date t , W_t , will simply be the dividends paid at date t , i.e. $W_t = w_m$. The wealth will be distributed among the investors proportionately according to their owned share of asset m . The distribution received by each agent, $W_{i,t}$, can be used to consume and re-invest. Following the discussion of Sciubba (1999), we assume that there is aggregate uncertainty so that $w_m \neq w_v$, for $m \neq v$.

There is a finite number of agents with *heterogeneous* temporal preferences in this economy, indexed by $i \in \{1, 2, \dots, I\}$. Each agent i has his subjective beliefs about the future sequence of the states. Each of these subjective beliefs is characterized by a probabilistic model, denoted by B^i . Since B^i may change over time, the time index t is added as B_t^i to make such a distinction. The agent’s objective is to maximize his life-time expected utility, and there are two decisions that are involved in this optimization problem. First, he has to choose a sequence of saving

rates starting from now to infinity, and second a sequence of portfolios to distribute his saving over M assets. Let us denote these two sequences of decisions by

$$\{\{\delta_{t+r}^i\}_{r=0}^\infty, \{\alpha_{t+r}^i\}_{r=0}^\infty\},$$

where δ_t^i is the saving rate at time t , and

$$\alpha_t^i = (\alpha_{1,t}^i, \alpha_{2,t}^i, \dots, \alpha_{M,t}^i)$$

is the portfolio comprising the M assets. The two sequences of decisions will be optimal and are denoted by $\{\delta_{t+r}^{i,*}\}_{r=0}^\infty$ and $\{\alpha_{t+r}^{i,*}\}_{r=0}^\infty$, if they are the solutions to the following optimization problem.

$$\max_{\{\{\delta_{t+r}^i\}_{r=0}^\infty, \{\alpha_{t+r}^i\}_{r=0}^\infty\}} E\left\{\sum_{r=0}^\infty (\beta^i)^r u^i(c_{t+r}^i) \mid B_t^i\right\} \quad (1)$$

subject to

$$c_{t+r}^i + \sum_{m=1}^M \alpha_{m,t+r}^i \cdot \delta_{t+r}^i \cdot W_{t+r-1}^i \leq W_{t+r-1}^i, \quad \forall r \geq 0, \quad (2)$$

$$\sum_{m=1}^M \alpha_{m,t+r}^i = 1, \quad \alpha_{m,t+r}^i \geq 0, \quad \forall r \geq 0. \quad (3)$$

In Equation (1), u^i is agent i 's temporal utility function, and β^i , also called the discount factor, reveals agent i 's time preference. The expectation $E(\cdot)$ is taken with respect to the most recent belief B_t^i . Equations (2) and (3) are the budget constraints. By combining constraint (3), constraint (2) can also be written as (4),

$$c_{t+r}^i \leq (1 - \delta_{t+r}^i) W_{t+r-1}^i, \quad (4)$$

where c_t^i denotes consumption. These budget constraints do not allow agents to consume or invest by borrowing.

Given the saving rate $\delta_t^{i,*}$, agent i will invest a total of $\delta_t^{i,*} \cdot W_{t-1}^i$ in the M assets according to the portfolio $\alpha_t^{i,*}$. In other words, the investment put into each asset m is $\alpha_{m,t}^{i,*} \cdot \delta_t^{i,*} \cdot W_{t-1}^i$. By dividing this investment by the market price of asset m at date t , $\rho_{m,t}$, one derives the share held by agent i of that asset, $q_{m,t}^i$.

$$q_{m,t}^i = \frac{\alpha_{m,t}^{i,*} \cdot \delta_t^{i,*} \cdot W_{t-1}^i}{\rho_{m,t}}, \quad m = 1, 2, \dots, M \quad (5)$$

The equilibrium price $\rho_{m,t}$ is determined by equating the demand for asset m to

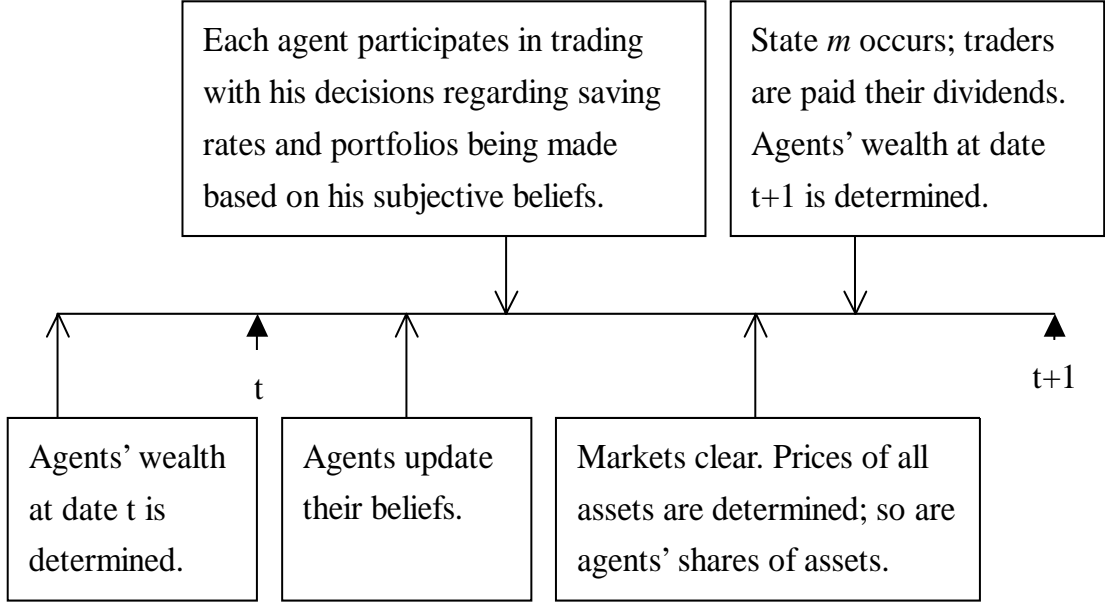


Fig. 1. The Time Line of a Market Day

the supply of asset m , i.e.,

$$\sum_{i=1}^I \frac{\alpha_{m,t}^{i,*} \cdot \delta_t^{i,*} \cdot W_{t-1}^i}{\rho_{m,t}} = 1, \quad m = 1, 2, \dots, M. \quad (6)$$

Rearranging Equation (6), one obtains the market equilibrium price of asset m :

$$\rho_{m,t} = \sum_{i=1}^I \alpha_{m,t}^{i,*} \cdot \delta_t^{i,*} \cdot W_{t-1}^i. \quad (7)$$

Agents' shares of assets will be determined accordingly by Equation (5). Afterwards, state m happens, and is made known to all agents at date t . The dividends w_m will be distributed among all stockholders of asset m in proportion to their shares, and their wealth will be determined accordingly as $W_t^i = q_{m,t}^i \cdot w_m$. The date moves to $t + 1$, and the process then repeats itself as shown in Figure 1.

3 The Agent-Based Multi-Asset Artificial Stock Market

The agent-based artificial stock markets share two generic features of the general ACE models, namely, *complex heterogeneity* and *bounded rationality*. In principle, agents can be heterogeneous in all of their characteristics, their decision rules, their beliefs and their preferences. However, these characteristics are not necessarily fixed as exogenously given. Agents' cognitive limits naturally bring about a trial-and-error process. Those characteristics may, therefore, be endogenously determined, and continuously change via learning and adaptation. The resultant

“biography” of agents has such potential to be very rich that Leigh Tesfatsion described it as the *artificial life*. The machinery used to grow the life of agents is termed *agent engineering*.

3.1 Agents’ Cognition

Like all agent-based computational economic models, we shall first begin 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 normal investors who try to maximize their lifetime discounted expected utility by appropriately choosing their *investment strategy*. The investment strategy is mainly composed of two parts, namely, *saving* and *portfolio*.

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

$$\{\{\delta_{t+r}^i\}_{r=0}^{\infty}, \{\alpha_{t+r}^i\}_{r=0}^{\infty}\},$$

where δ_t^i is the saving rate at time t , and α_t^i 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 Equations 1, 2, 3).

The discrete-time stochastic optimization problem defined by Equations (1), (2), and (3) may be analytically solvable when one considers some specific types of utility functions with some other necessary simplifications. However, since the purpose of this paper is to examine the relevance of risk preference to survivability, we place little in the way of restrictions on the types of risk preference. In fact, in this paper we even allow for the risk preference to be randomly generated, as long as it is well-behaved. This causes the optimization problem, when generally posed, to be difficult to solve analytically, not to mention the further complications arising from their beliefs or the conditional expectations.⁷ Therefore, we assume that all agents in our model are *computational*. They cope with the optimization problem with a numerical approximation method, and the specific numerical method used in this paper is the *genetic algorithm*.

⁷ For the purpose of providing an existential proof, there is no such need to solve this generally-posed problem. Instead, it is sufficient to assume the existence or the emergence of these optimizing agents (Blume and Easley, 1992; Sandroni, 2000). However, for us, even though the problem can be solved analytically, the agents may still be not optimizing, since they may misperceive the future prices of assets $m \{\rho_{m,t+r}\}_{r=0}^{\infty}$ and that may cause their *ex-ante* (planned) shares of assets, $\{q_{t+r}^i\}_{r=0}^{\infty}$, to be different from their *ex-post* (realized) ones. Spear (1989) has shown that for markets composed of complex heterogeneous agents, the rational expectations equilibria may not even be computable.

Over the last few years, the *genetic algorithm* has been the most active tool 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. In this paper, we use the genetic algorithm to evolve both 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 (1). First, the cognitive limit of investors and the resultant adaptive behavior free them from an infinite-horizon stochastic optimization problem, as in Equation (1). Instead, due to their limited perception of the future, the problem effectively posed to them is the following:

$$\max_{\{\{\delta_{t+h}\}_{h=0}^{H-1}, \{\alpha_{t+h}\}_{h=0}^{H-1}\}} E\left\{ \sum_{h=0}^{H-1} (\beta^i)^h u^i(c_{t+h}^i) \mid B_t^i \right\} \quad (8)$$

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_t^i , where B_t^i 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_t^i can be just the *subjective probability function*, i.e.

$$B_t^i = (b_{1,t}^i, \dots, b_{M,t}^i), \quad (9)$$

where $b_{m,t}^i$ is investor i 's subjective probability of the occurrence of the state m in any of the next H periods. In a more general setting, B_t^i can be a *high-order Markov process*. With the replacement (8), 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_t^i is *adaptive*.

Furthermore, we assume that investors will *continuously* adapt their investment strategies according to the *sliding window* shown in Figure 2. At each point 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 simplifi-




Fig. 2. A Sliding-Window Perception of the Investors

cations of the optimization problem (1) – (3). 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 ,

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

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.

$$\delta_t^i = \delta_{t+1}^i = \delta_{t+2}^i = \dots \delta_{t+H-1}^i, \quad (11)$$

$$\alpha_t^i = \alpha_{t+1}^i = \alpha_{t+2}^i = \dots \alpha_{t+H-1}^i. \quad (12)$$

With these two simplifications, we replace the original optimization problem, (1) – (3), that is presented to the infinitely-smart investor, with a modified version which is suitable for a boundedly-rational investor.

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

subject to

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

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

$$c_{t+h}^i = (1 - \delta_t^i) W_{t+h-1}^i, \quad \forall h \in \{0, H-1\}. \quad (16)$$

3.2 Autonomous Agents

One of the mainstays of agent-based computational economics is *autonomous agents* (Tesfatsion, 2001). The idea of autonomous agents was initially presented in Holland and Miller (1991). Briefly, these agents are able to learn and to adapt to the changing environment without too much external intervention, say, from the model designer. Their behavior is very much endogenously determined by the environment with which they are interacting. Accordingly, sometimes it can be very difficult to trace and to predict, and is known as *emergent behavior*.

In this paper, we follow what was initiated in Holland and Miller (1991), and equip our agents with the genetic algorithm to cope with the finite-horizon stochastic dynamic optimization problem, (13) – (16). The GA is applied here at two different levels, a high level (learning level) and a low level (optimization level). First, at the high level, it is applied as a *belief-updating scheme*. This is about the B_t^i appearing in (13). Agents start with some initial beliefs of state uncertainty, which are basically characterized by parametric models, say, Markov processes. However, agents do not necessarily confine themselves to just stationary Markov processes. Actually, they can never be sure whether the underlying process will change over time. So, they stay alert to that possibility, and keep on trying different Markov processes with different time frames (time horizons). Specifically, each belief can be described as “a k th order Markov process that appeared over the last d days and may continue”. These two parameters can be represented by a binary string, and a canonical GA is applied to evolve a population of these two parameters with a set of standard genetic operators. Details are given in Section A.2.

Once the belief is determined, the low-level GA is applied to solve the stochastic dynamic optimization problem defined in (13) – (16). Basically, we use Monte Carlo simulation to generate many possible ensembles consistent with the given belief and use them to evaluate a population of investment plans composed of a saving rate and a portfolio. GA is then applied to evolve this population of candidates. Details are given in Section A.1.

In sum, the high-level GA finds an appropriate belief, and under that belief the low-level GA searches for the best decisions in relation to savings and portfolios. This style of adaptive design combines *learning how to forecast* with *learning how to optimize*, a distinction made in Bullard and Duffy (1999). These two levels of GA do not repeat with the same frequency. As a matter of fact, the belief-updating scheme is somewhat slow, whereas the numerical optimization scheme is more frequent. Intuitively, changing our belief of the meta-level of the world tends to be slower and less frequent than just fine-tuning or updating some parameters associated with a given structure. In this sense, the idea of *incremental learning* is also applied to our design of autonomous agents.

3.3 The Behavior of CAPM Believers

Investors whose behavior is governed by the procedure described in Sections A.1 and A.2 above are autonomous, which may not come to any stereotype familiar to us. To make sense of the evolution going on in the artificial markets, it would be useful to include some familiar types of investors as well, for example, those investors, no matter what happens, who just follow a guideline which is exogenously given. These investors are not autonomous. We shall call them *formula investors*. The GA procedure will not be applied to these formula investors. Instead, their behavior is prespecified by a formula. As an illustration of the idea of formula investors, we follow Sciubba (1999) to introduce *CAPM believers* into the market.

The CAPM believers are investors who base their portfolio rule upon the well-known *capital asset pricing model (CAPM)*. In the spirit of the CAPM, they first find out the *market portfolio* and the *risk-free portfolio*. Then, according to their risk preference (degree of risk aversion), they choose a weighted combination of the two. We index CAPM traders by means of κ , and let γ^κ be the associated risk aversion coefficient, which is randomly determined by the uniform distribution $U(0,1)$. At date t , each investor κ invests in asset m a portion $\alpha_{m,t}^{CAPM(\kappa)}$ of his savings such that

$$\alpha_{m,t}^{CAPM(\kappa)} = \gamma^\kappa \alpha_{m,t}^F + (1 - \gamma^\kappa) \alpha_{m,t}^M, \quad m = 1, \dots, M, \quad (17)$$

where $\alpha_{m,t}^F \equiv \frac{\hat{\rho}_{m,t}/w_m}{\sum_{m=1}^M \hat{\rho}_{m,t}/w_m}$ and $\alpha_{m,t}^M \equiv \frac{\hat{\rho}_{m,t}}{\sum_{m=1}^M \hat{\rho}_{m,t}}$.⁸

This way of defining the CAPM investment strategy is proposed by Sciubba (1999). It is straightforward to show that the vector α_t^M is in effect the weighted average of all market participants' portfolios in the previous period. So, it is intuitively consistent with the idea of the market portfolio. Furthermore, it can be shown that the vector α_t^F is a weighted average of the price-earnings ratio, and with this portfolio the agent can expect to earn the same rate of return regardless of the occurrence of the state. As a result, it can be taken so as to approximate the idea of the risk-free rate.

Similarly, one can come up with a respective saving rate for the CAPM investor as follows:

$$\delta_t^{CAPM(\kappa)} = \gamma^\kappa \delta_t^F + (1 - \gamma^\kappa) \delta_t^M \quad (18)$$

where $\delta_t^F \equiv 0$ and $\delta_{m,t}^M \equiv \sum_{i=1}^I \left(\frac{W_{t-2}^i}{\sum W_{t-2}^i} \right) \cdot \delta_{t-1}^i$.

The CAPM believers' adherence to the CAPM portfolio rule (17) is not affected during the entire course of evolution.

⁸ For simplicity, we may assume that they have static expectations as other autonomous agents do, i.e. $\hat{\rho}_{m,t} = \rho_{m,t-1}$, $\forall m = 1, \dots, M$.

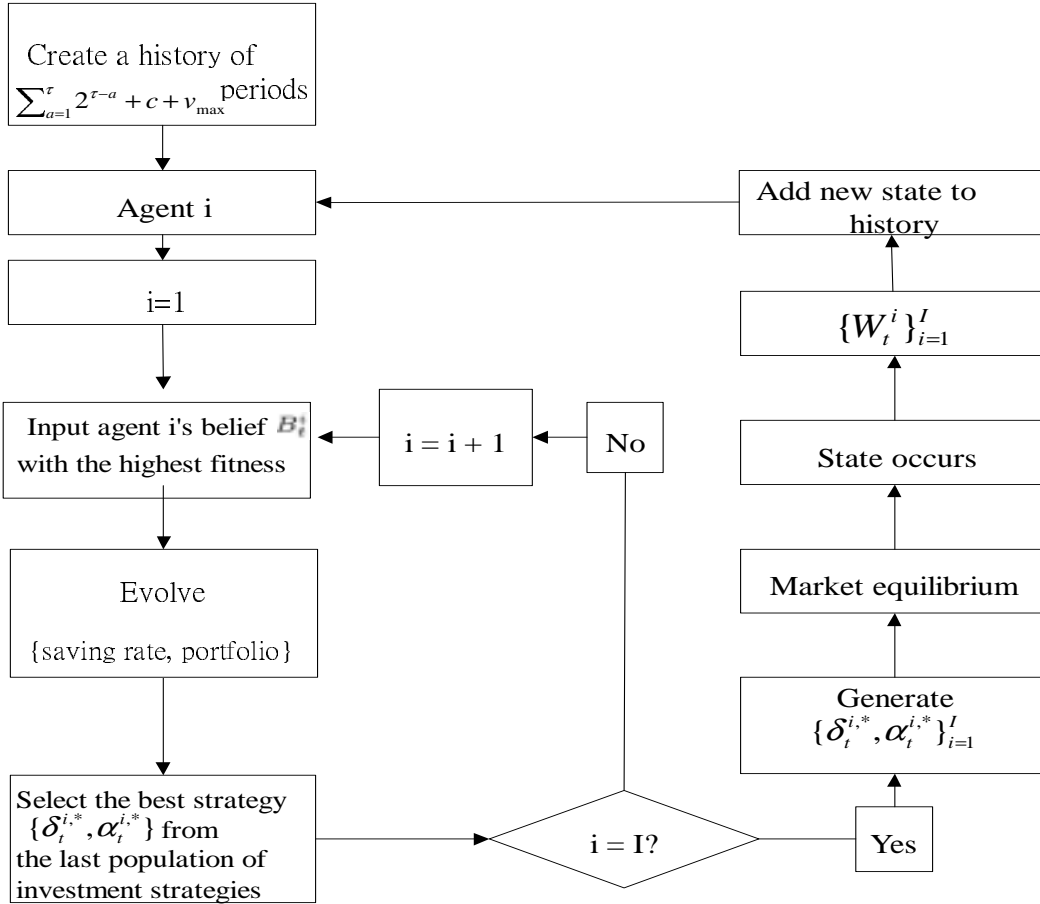


Fig. 3. A Summary of Agent-Based Artificial Stock Markets

3.4 Summary

Figure 3 is a summary of the agent-based artificial stock market.

4 Experimental Designs

Two experiments are conducted in this agent-based artificial stock market. In the first experiment, the autonomous agents are identical in all respects except for their *preferences over risk*. The purpose of this design is to see whether the survivability of investors has anything to do with their preferences. In the second experiment, we further distinguish each type of agent by their *belief formation processes* which, under some circumstances, are pertinent to forecasting accuracy. Obviously, the second design is to examine the relevance of forecasting accuracy to survivability. We believe that these two designs together will contribute to the resolution of the debate of the irrelevance of risk preference or the dominance of forecasting accuracy in a rich empirical context.

4.1 Market and Participants

In both experiments, the market is composed of 40 agents ($I = 40$). Five out of the 40 are CAPM believers. Since their portfolio and saving decisions are determined by (17) and (18) and are irrelevant to their preferences over risk, their preferences will not be specified here. The remaining 35 agents are all autonomous. Their behavior is mainly driven by what has been detailed in Sections A.1 and A.2. For these agents, the specification of their preferences is required, and they are detailed in Table 1.

In total, we consider seven types of autonomous agents, and each type is assigned to five agents. Type one has the logarithmic utility function. We are very much interested in knowing whether this type of agent has any advantage over others in terms of the long-run wealth share. As to types two to six, they are also frequently used in economic analysis.⁹ Among them, type four is the well-known *CARA* (constant absolute risk aversion) utility function. In addition to these six familiar types of utility functions, we also consider any arbitrary utility function. By using Taylor expansion, an arbitrary analytical utility function can be approximated by a finite-order polynomial function. Here, we consider the approximation only up to the sixth order.

Notice that types 3 to 7 refer to a class of parametric utility functions. Parameters of these types of utility functions, namely, $\alpha_1, \dots, \alpha_4, \beta_1, \dots, \beta_3, \gamma_3$ and a_0, a_1, \dots, a_6 , can in principle be randomly or manually generated as long as they satisfy the regular first- and second-order conditions: $u' > 0$ and $u'' < 0$. Since each type of utility function is assigned to five agents, parameter values are generated for each agent of each type separately. So, agents of type 3 may have different values of (α_1, β_1) , agents of type 4 may have different values of (α_2, β_2) , and so on and so forth.

There are 5 assets available in the market ($M = 5$), corresponding to 5 states. Asset m pays dividends $6 - m$ ($m = 1, 2, \dots, 5$). Two stochastic processes are considered in the experiments, namely, *iid* and the *first-order Markov*. Each is employed for one half of the total number of runs. Parameters of these two stochastic processes are also randomly generated in such a way that the axioms of the probability function are satisfied.

4.2 Parameters related to Autonomous Agents

At each point in time, agents have a perception of a time horizon with length $H = 25$. To solve the optimization problem (8), agents simulate 5 25-horizon en-

⁹ See, for example, Huang and Litzenberger (1988), pp. 27-33.

Table 1
Types of the Utility Function $u(c)$

	Utility Type	RRA
Type 1	$u(c) = \log(c)$	1
Type 2	$u(c) = \sqrt{c}$	0.5
Type 3	$u(c) = \alpha_1 + \beta_1 c$	0
Type 4	$u(c) = \frac{\alpha_2}{\beta_2} \exp\{\beta_2 c\}$	$-\beta_2 c$
Type 5	$u(c) = \frac{1}{(\gamma_3+1)\beta_3}(\alpha_3 + \beta_3 c)^{\gamma_3+1}$	$-\frac{\beta_3 \gamma_3}{\frac{\alpha_3}{c} + \beta_3}$
Type 6	$u(c) = c - \frac{\alpha_4}{2} c^2$	$\frac{\alpha_4}{\frac{1}{c} - \alpha_4}$
Type 7	$u(c) = a_0 + \sum_{i=1}^6 a_i c^i$	$-\frac{2a_2 c + 6a_3 c^2 + 12a_4 c^3 + 20a_5 c^4 + 30a_6 c^5}{a_1 + 2a_2 c + 3a_3 c^2 + 4a_4 c^3 + 5a_5 c^4 + 6a_6 c^5}$

sembles ($L = 5$) of the states based on their belief in order to evaluate the fitness of their investment strategies (A.4).

The design for the agents' adaptation is composed of two parts. For the low-level evolution, the architecture is a population GA. The population size N (number of investment strategies) for each agent is 100. The genetic parameters applied to evolving this population are as follows. The crossover rate (p_{cross}^l) is set to be 1, while the mutation rate (p_{mutate}^h) is set to be 0.03. Tournament selection with a tournament size of 4 is applied. The number of generations that the low-level GA runs in one period, G , is set to be 50.

As to the high-level evolution, the architecture is also a population GA, and the population size J (number of beliefs) maintained by each agent is one hundred. The crossover rate and the mutation rate are the same as those of the low-level evolution, as is the tournament size used for the tournament selection. The belief set will be renewed after every 2 periods ($\Delta = 2$).

Table 2 provides a summary of the design. Two experiments are conducted based on this common design. The two experiments mainly differ in the parameter v , which is the size of the data used to validate the model. In Experiment 1, our focus is on the role of the utility function. As a result, the v is fixed for all autonomous agents, which is 100. Nonetheless, since the stochastic process simulated in this market is *stationary*, it is expected that a larger v will help validate the model and enhance the forecasting accuracy. Therefore, to see how significant the forecasting accuracy can be (Sandroni's main argument), in Experiment 2 we let autonomous agents have different values of v , starting from the very small one, 10, increasing it to 15, 25, 50, and finally to the largest one 100. To not mix the results of the two experiments, we let these five values of v be evenly distributed among each type of agent, i.e. one for each agent of each type. With this distinction, the two experiments together can help us see better which factor is more important for the survivability of agents, namely, the *preference* or *forecasting accuracy*.

Table 2
Experimental Design

Market and Participants	
Number of market participants (I)	40
Number of types of agents	8
Number of each type of agent	5
Number of assets (states) (M)	5
Dividends paid by asset m	$6 - m$
Stochastic processes	iid or 1st-order Markov
Number of market periods (T)	100
Discount rate (β)	0.45
Autonomous Agents	
Agents' perception of the time horizon (H)	25
Number of ensembles (L)	5
Population size (number of strategies) (N)	100
Number of generations (G)	50
Population size (number of beliefs) (J)	100
Frequency of running GA on the belief set (Δ)	2
Crossover rate (p_c^l, p_c^h)	1
Mutation rate (p_m^l, p_m^h)	0.05, 0.03
Tournament size	4
Number of bits for beliefs ($\tau_1 + \tau_2$)	10

The eight types refer to the seven types of autonomous agents whose utility functions are specified in Table 1 plus the CAPM believers.

Both experiments are run 100 times, and each run lasts for 100 market periods ($T=100$). The simulation is conducted through the software *AIE-ASM Version 5.0*. The software is written with *Delphi, Version 6.0*. Usually, it takes 3 hours for a single run on a Pentium III 1000 with 256 MB RAM personal computer.

The γ values exogenously given to 5 CAPM believers are 0.1, 0.2, 0.3, 0.4 and 0.5, respectively.

5 Experimental Results

In each single run, we generate a series of artificial data. At the micro level, it includes the dynamics of agents' beliefs, investment behavior, and the associated wealth

$$\{B_t^{i,*}, \delta_t^i, \alpha_t^i, W_t^i\}_{t=1}^{100}, \quad i = 1, \dots, 40.$$

At the aggregate level, we observe the asset price dynamics

$$\{\rho_{m,t}\}_{t=1}^{100}, \quad m = 1, \dots, 5.$$

Since the main concern of this paper is with which types of agent survive, our focus is on the *wealth share dynamics*. In addition, in order to understand what makes surviving agents survive, agents' belief dynamics and investment behavior also attract our attention. So, the result presented in this section will be pretty much based on the micro-level data. Little will be said of the price dynamics.¹⁰

5.1 Experiment 1

5.1.1 Wealth Share Dynamics

Figure 4 shows the wealth-share dynamics of the eight types of investors. Notice that each line is based on the average of 100 simulations. The results clearly demonstrate the strong dominance of the type-one investors. While in some cases type-two investors are still hanging in there till the end of the 100-period simulation, their shares, in all of our 100 simulation runs, are smaller than those of the type-one investors, and many of them are declining toward zero.¹¹ It is also interesting to notice that the type-one investor has a constant relative risk aversion coefficient that is one. Therefore, our findings can lend support to Blume and Easley's main argument: *the market selects those investors whose coefficient of relative risk aversion is nearly one*.¹² In fact, if we consider the family of the CRRA utility functions, namely,

¹⁰ Some preliminary time series analysis of price dynamics can be found in Chen and Huang (2004).

¹¹ One may suspect that if the number of iterations (T) is long enough, say $T=500$, then the type-1 agents are the only type of survivors. Actually, in a separate experiment, we have found that this is indeed the case. (Chen and Huang, 2004)

¹² See Blume and Easley (1992), Theorem 5.4, pp. 23-24. The italics shown in the main text are not the exact quotation of that theorem, which was originally based on controlling saving rates. Since saving rates are treated endogenously in our paper, our finding suggests that the theorem can still be true even though the assumption regarding saving rates is relaxed. Theorem 5.4 also rests upon forecasting accuracy, which we will address later.

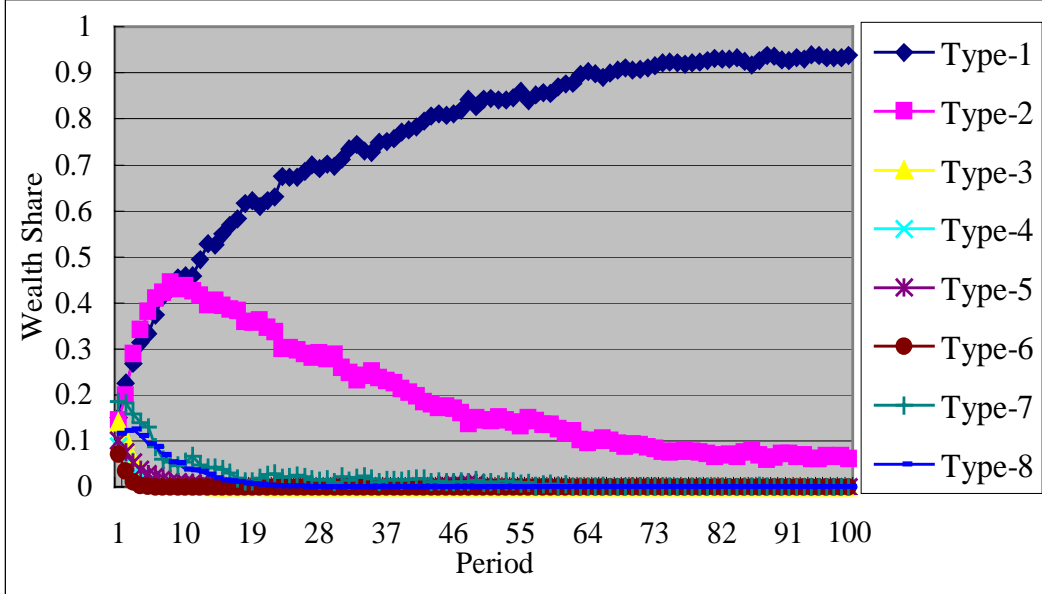


Fig. 4. Time Series Plot of the Wealth Share of Eight Types of Investors: Experiment 1

$$u(c) = \begin{cases} c^\rho/\rho, & \text{if } -\infty < \rho < \infty \text{ and } \rho \neq 0, \\ \ln^c, & \text{if } \rho = 0, \end{cases} \quad (19)$$

Type-2 investors $u(c) = \sqrt{c} = c^{\frac{1}{2}}$ also has a constant RRA coefficient of 0.5 ($1-\rho$), but it is not close enough to one.¹³

5.1.2 Forecasting Accuracy

Sandroni, however, considered *forecasting accuracy* to be the sole important factor in determining who may survive. Since the design of Experiment one equip all autonomous agents, regardless of their types, with the same learning scheme, namely, genetic algorithms, intuitively it would be hard to attribute the survival of type-one agents to their proficiency in forecasting. Nevertheless, in this section, we shall take a closer look at the forecasting performance of different types of agents.

The *Kolmogorov-Smirnov* statistic (KS statistic hereafter) is chosen to measure the forecasting accuracy of agents. The KS statistic is a metric, or more precisely the *sup norm*, for two distribution functions. Formally, let F and G be two distri-

¹³ In a separate study (Chen and Huang, 2004), we simulated the cases with different ρ s from 0 to 0.9 with an increment of 0.1. It was found that the agents' share of wealth positively related to their relative risk coefficients, $1 - \rho$. The wealth share dynamics of agents with a ρ of 0.1 is very close to that of type-1 agents considered here. Both survive to the end of the 100-period simulations.

Table 3
Performance Measurements

Forecasting Accuracy (K-S statistics)						
Type 1	Type 2	Type 3	Type 4	Type 5	Type 6	Type 7
0.05219	0.05220	0.05233	0.05228	0.05222	0.05226	0.05220
Portfolio Performance						
<i>Mean Rate of Return</i>						
2.38024	2.83639	2.80554	2.79910	2.80125	2.92166	2.78132
<i>Standard Deviation of Return</i>						
7.28998	8.10897	8.33604	8.91320	9.56133	10.91684	8.12857
<i>Sharpe Ratio</i>						
0.32651	0.34978	0.33656	0.31404	0.29298	0.26763	0.34217

tribution functions

$$KS(F, G) = \sup_x | F(x) - G(x) | . \quad (20)$$

Using the KS statistic, we can measure the forecasting error by the difference between the true distribution function and the subjective distribution perceived by agents. The results are shown in Table 3 (the top panel). The statistical test cannot reject the null hypothesis that the forecasting accuracy among different types of autonomous traders is equally good. As already mentioned, this test result is not surprising given the fact that all autonomous agents are supplied with the same adaptive scheme to update their beliefs. Therefore, forecasting accuracy, at least, is not the sole important factor in the determination of survivability.

5.1.3 Saving Rates

Since forecasting accuracy cannot necessarily guarantee agents' survival, there are only two decision variables left for us to see the uniqueness of type-1 agents, namely *saving* and *portfolio*. Notice that in Blume and Easley (1992), the saving rate is exogenously given, and it is found that the saving rate can play an important role in determining who survives.¹⁴ However, its impact becomes *implicit* when the saving decision is endogenized (Sandroni, 2000).¹⁵ It is therefore useful to

¹⁴ For example, see Blume and Easley (1992), proposition 3.2, p.16.

¹⁵ In his illustrating example on p. 1311, Sandroni (2000) compared two agents with different preferences: the one with the square-root utility function had the correct prediction, whereas the other with the log utility function did not have the correct prediction. On p. 1313, he then showed that the former with a higher saving rate drove out the latter with a lower saving rate, and the difference in their saving rates was endogenously generated as a part of the equilibrium.




Fig. 5. The Box-and-Whisker Plot of the Lifetime Saving Rates of the Seven Types of Agents (Experiment 1)

examine the saving behavior of different types of agents, if one desire to know the survival of type-1 agents.

Figure 5 is a box-whisker plot of the saving rates among the seven types of agents. Each plot shows the life-time distribution of the saving rate δ_t associated with a specific type of agent. To generate each plot, we first take an average of the saving rate of the five agents of the same type. This is done period by period. Then we have a time series of the saving rate δ_t ($t = 1, 2, \dots, 100$) for each type of agent. To have an idea of the distribution (dispersion) of the saving behavior, the minimum, the first quartile, the medium, the third quartile, and the maximum of $\{\delta_t\}$ are recorded for each type of agent. Furthermore, we then derive sample statistics for these order statistics by taking an average over the entire 100 simulation runs. These statistics are depicted in the form of a box-whisker plot for each type of agent as shown in Figure 5.

The line appearing in the middle of the box indicates the median saving rates for a specific type of agent. While higher saving rates, as Blume and Easley suggested, will place agents in an advantageous position to survive, we find that the saving rate of type-1 agents is in fact the lowest among all seven types of agents. This is evidenced by the lowest median of the type-1 agents. What, however, makes type-1 agents unique is their *very stable saving behavior*. This is revealed by comparing the boxes and whiskers of the plots. As opposed to other types of agents, type-1 agents obviously have a very narrow box with a very short whisker, which features a *very stable saving behavior*. Not only is it stable, but it is also stable around 0.45, which is exactly the discount rate β set in Table 2.

From an analytical viewpoint, Blume and Easley (1992) have already showed that $\delta_t = \beta$ for log-utility agents. In other words, the optimal saving rate for type-1 agents is just a constant and is independent of beliefs, wealth and asset prices. Our genetic algorithm (the low-level GA) just confirms this property *numerically*. We think that this property has an important implication for the survival of type-1 agents, and our reasons for this follow.

In general, the saving decision is made jointly with the portfolio decision, which means that when agents are equipped with a GA, the low-level GA actually works with the high-level GA. While an agent uses the low-level GA to make the saving and portfolio decision, the quality of that decision also depends on the belief (forecasting accuracy) supplied by the operation of the high-level GA. Hence, any imperfection in the high-level GA may compound the imperfection of the low-level GA in the usual sense of error propagation. Nevertheless, since for type-1 agents the saving decision is independent of their beliefs and hence forecasting accuracy, it separates the performance of the low-level GA from that of the high-level GA. It behaves like this. The low-level GA first learns that β is the only relevant factor for the saving decision, and simply works with β . Furthermore, since β is a noise-less constant, the saving decision associated with the low-level GA is well-grounded. In the end, it helps type-1 agents come up with a quality decision on the saving rate, that is almost exact.

The same story above unfortunately does not apply to other types of agents. In general, their saving decisions are not independent of their beliefs. Even though these agents are equipped with the same kind of GA used by type-1 agents, their low-level GA is operated under the belief determined by the high-level GA, which may suffer from some degree of inaccuracy all the time. This will in turn have an adverse effect on the quality of the saving decision. From what we have seen in Figure 5, this happens in terms of the appearance of the large fluctuation in saving rates. In addition to the unstable saving behavior, type-3 to type-7 agents have suffered further from the extremely low *down-side* saving rates, which may contribute to the fast decline in their wealth share. Type-2 agents do not share this feature of down-side saving rates. Obviously they perform better, but their unstable saving behavior eventually causes a threat to them.¹⁶

It is worth noting that the significance of a stable saving behavior was not a focus of Blume and Easley's original analysis, when they treated the saving rate as an exogenously-given constant. It is also hard to address this in Sandroni's dy-

¹⁶ In this general equilibrium setting where almost everything is endogenously determined, error can propagate in a quite complex fashion. For example, when the saving rate is incorrectly determined, it may impact the asset holding of the agents, and further impact the associated dividends received and then wealth. The resultant wealth may distort the saving decision further since the latter in general depends on the former. In addition to wealth, error can also propagate through the endogenously generated price, e.g., a very fluctuating price which may make the saving decision even more difficult.

dynamic equilibrium framework, where the learning dynamics is absent.¹⁷ Hence, our agent-based simulation provides a new idea to ponder over as to why preference matters or, more specifically, why log-utility agents can drive out all other agents. The evidence brings our attention to the quality of saving behavior and its contributing factors in the light of learning dynamics.

Stable saving behavior is not new in economic analysis. It has already drawn the attention of economists in the recent studies on the lock-up savings.¹⁸ However, the implied stable saving behavior observed in our simulation provides a different insight.

5.1.4 Portfolio Performance

The second decision is regarding *portfolio*. The log-utility (type-1) agents, as described by Blume and Easley (1992), are the kind of agents who maximize their expected growth rate of wealth, also known as the *Kelly criterion*. This observation naturally raises a question: would type-1 agents survive because their portfolio performance is superior to that of other types of agents? To answer this question, we can start with two different performance measurements: one is *ex ante*, and the other is *ex post*. The former evaluates the agents' investment based on its (probabilistic) expected value, while the latter evaluates it on the basis of its realized value. In this paper, we consider the ex-post approach to be more pertinent. There are two reasons for us to think so.

First, agents in the ACE model are not optimizing agents; instead, they are adaptive agents. The adaptation scheme, be it GA or not, is generally driven only by realized returns rather than by the expected returns. It is well known that such a discrepancy can deviate agents' behavior from that of expected-utility maximizers (Lettau, 1997). This leads us to the second point: in the real world, what is used to evaluate or rank the performance of mutual fund managers certainly is the historical returns and not the hypothetical expected returns. Having said that, we suggest calculating the rate of return as follows. We first calculate $r_{m,t}^i$, the realized rate of return of agent i on the investment in the asset m at time t , as

$$r_{m,t}^i = \frac{w_m \left(\frac{\delta^i W_{t-1}^i \alpha_{m,t}^i}{\rho_{m,t}} \right) - \delta^i W_{t-1}^i \alpha_{m,t}^i}{\delta^i W_{t-1}^i \alpha_{m,t}^i} = \frac{w_m}{\rho_{m,t}} - 1. \quad (21)$$

Weighting $r_{m,t}^i$ over all assets ($m = 1, \dots, M$) by its associated portfolio, one then

¹⁷ However, one important observation in Sandroni's analysis, when he endogenized the saving decision, is that the consequence of forecasting accuracy should not just be limited to the portfolio decision, but also the saving decision. This is by and large consistent with our analysis above.

¹⁸ See, for example, Laibson (1998).

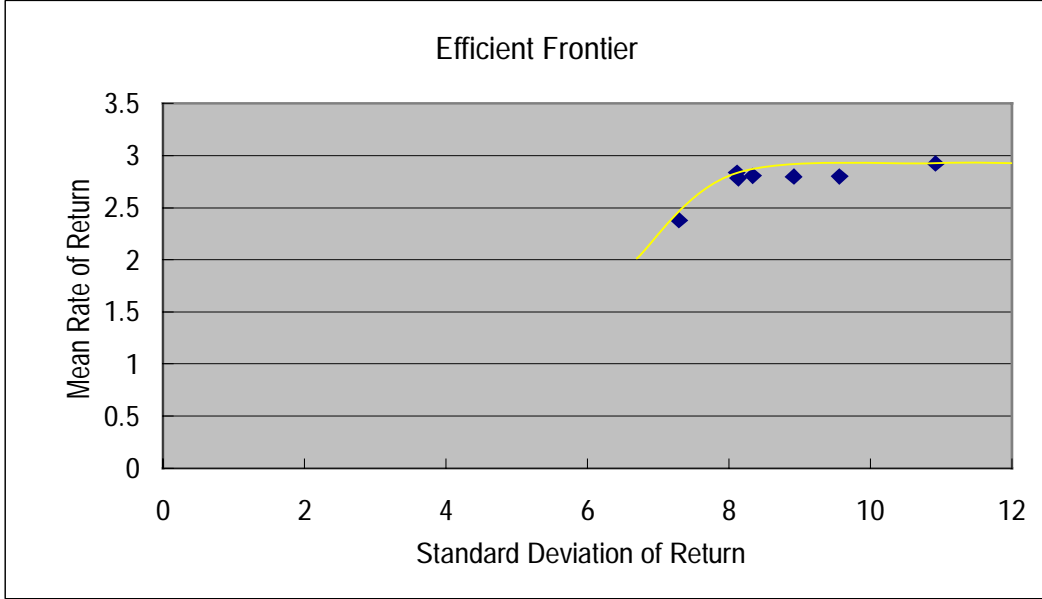


Fig. 6. The Efficient Frontier

derives the rate of return of agent i for period t

$$r_t^i = \sum_{m=1}^M \alpha_{m,t}^i r_{m,t}^i. \quad (22)$$

Let \bar{r}^i and $\text{var}(r)^i$ denote the *sample mean* and the respective *sample variance* of the series $\{r_t^i\}$. Table 3 shows these two statistics, which are averaged over the entire 100 simulation runs and are further averaged over the five agents of the identical type.

From Table 3, we learn that the type-1 agents do not survive because of their superior performance in the expected rate of return. As a matter of fact, among all the seven types of agents, it is they that have the lowest rate of return (2.38), which is different from what one may expect from the Kelly criterion. Nonetheless, the column “standard deviation of return” indicates that these agents are under different exposure to risk, and it is the type-1 agents who are exposed to the lowest risk. This result is not totally unanticipated given the fact that the type-1 agents are the most risk-averse.¹⁹ Motivated by this finding, we go further to examine the *risk-adjusted return*, also known as the *Sharpe ratio*, and we find that type-1 agents do not perform particularly well in terms of the Sharpe ratio.

¹⁹ This can be seen from the relative risk aversion (RRA) coefficient revealed from Table 1. It is not straightforward to compare the RRA coefficient of the type-1 agents with those of the type-4 to type-7 agents, which are not constant and change with consumption. However, in the particular case when their consumption gets closer to zero, their RRA coefficients also converge to 0. Since this case indeed applies quite frequently to our simulations, effectively speaking their RRA coefficients are also lower than that of the type-1 agents.

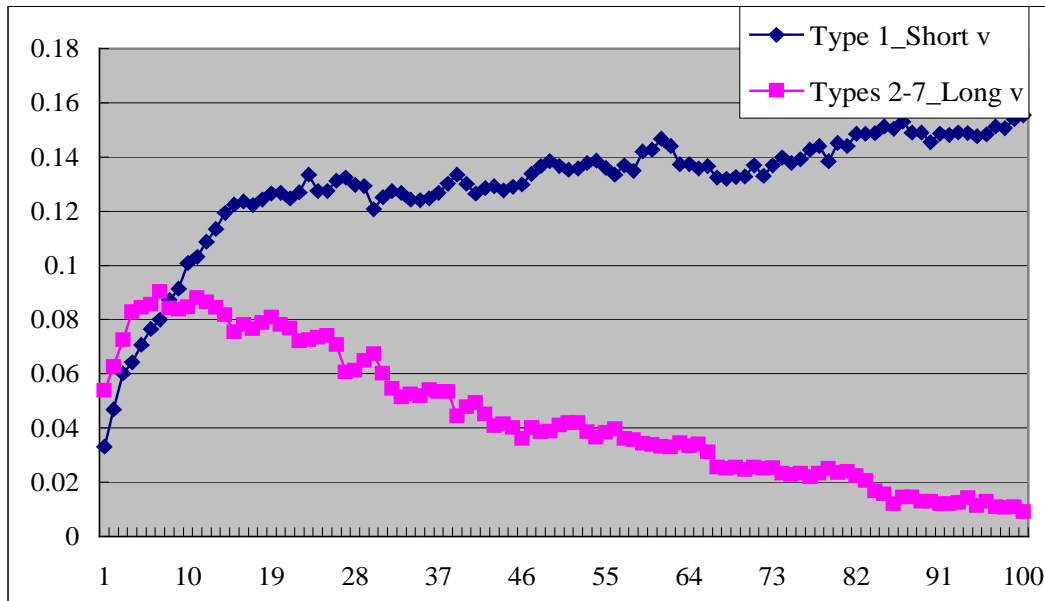


Fig. 7. Time Series Plot of the Wealth Share Per Capita: Experiment 2 (Type-1 vs Others)

Especially, one may suppose that every investor whose performance is situated at the *efficient frontier* has an equal chance to survive.²⁰ Therefore, we see no particular reason to attribute the survival of the type-1 agents to their portfolio performance.

5.2 Experiment 2

In Section 5.1.2, we already saw that agents with identical capabilities of forecasting do not survive equally well. This result is already evidence that forecasting accuracy is not the primary force in the determination of survivability. Instead, preference plays a dominating role. However, since all agents forecast equally well (Table 3), the result obtained in Experiment 1 does not lend strong support to the evidence that shows that forecasting accuracy does not matter at all. Therefore, to see whether we can consolidate the argument that forecasting accuracy does not matter, a different experimental design is proposed in this section.

Experiment 2 assigns different values of v , i.e. different horizons of the validation sample (see Equation (A.10)). Since the underlying dividends series is stationary, doing this will normally lead those agents using a long validation horizon to outperform those agents using a short validation horizon.²¹ Now, given the dif-

²⁰ To see this, the risk-return plot is drawn in Figure 6. The continuous frontier line is constructed by smoothly connecting the three points on the frontier. The three points on the frontier correspond to type-1, type-2 and type-6 agents. While the other four types of agents do not lie exactly on the frontier, they are not far away from it.

²¹ This result is confirmed by the K-S statistic. The K-S statistic of the five validation




Fig. 8. Time Series Plot of the Wealth Share Per Capita: Experiment 2 (Across Validation Horizons)

ference in their forecasting accuracy, it is time to ask whether agents, regardless of their preference types, with superior forecasting accuracy can survive better as opposed to those agents with inferior forecasting accuracy, which is also the main issue addressed in Sandroni's proposition 3. What is particularly relevant to our concern is whether agents of Types 2 to 7 who have long validation horizons can drive out those Type-1 agents who have short validation horizons.

In Figure 7, we classify agents with different preferences into two groups, namely, agents of type one alone, and agents of types 2 to 7 together, and plot the time series of the wealth share dynamics of these two different groups of agents. Additional restrictions are added to these two groups. For type-1 agents, only those whose validation horizons are short (to be specific, $v = 10, 15$) are included, and the respective time series plot is denoted by diamonds. For the remaining two groups of agents, we only consider those with long horizons ($v = 50, 100$), and the time series plots are denoted by squares. These contrasting plots help elucidate the effect of forecasting accuracy on survival.

From Figure 7, we can see that, even though the validation horizon is lengthened to enhance the forecasting accuracy of agents of types 2 to 7, this group of agents do not perform well: their wealth share per capita continues to decline toward nil. Therefore, it is clear that forecasting accuracy does not help them much to survive. Nonetheless, if we restrict our attention to agents with different validation

schemes, $v=10, 15, 25, 50$ and 100 are $0.0757, 0.0695, 0.0623, 0.0551,$ and 0.0497 , respectively. The ANOVA analysis shows that they are significantly different with an F statistic of 26.41 .

horizons, then the significance of forecasting accuracy is revealed. Figure 8 shows the wealth dynamics per capita by polling agents with the same validation horizon together. While the wealth dynamics fluctuates quite severely, the general tendency indicates that, regardless of their risk preference, agents with better forecasting accuracy tend to have higher wealth shares than agents whose forecasting accuracy is worse. Nevertheless, a lower wealth share is certainly not good enough to satisfy the market selection hypothesis, which actually predicts that those agents with worse forecasting accuracy will *vanish*.

Therefore, forecasting accuracy, as our conventional wisdom may suggest, does matter for the prosperity of agents. It is significant when agents share the same risk preference. Nevertheless, when agents have heterogeneous risk preference, its importance is only secondary, and the effect of risk preference may dominate.²²

6 Concluding Remarks

Normally, if we deviate from the original theoretical assumptions, for example, by introducing bounded-rational behavior, we may expect some different results. This is not surprising, since the learning literature over the past decade has accumulated much evidence of this kind. It would be interesting, though, if the new result is *strikingly* different from the previous findings. We believe that to be the case here. Earlier theoretical studies have already shown the irrelevance of risk preference to survivability. They have also shown that what matters is forecasting accuracy. After introducing bounded-rational behavior, we have almost the opposite: risk preference matters and it is even more important than forecasting accuracy.

Why does risk preference matter? The paper examines three interdependent possibilities: forecasting accuracy, portfolio and saving among different types of bounded-rational agents. While being bounded rational, all of these agents are at least potentially equally smart in the sense that they are all equipped with the same adaptive search scheme, namely, the genetic algorithm. The neutrality of GA, in that GA does not make one type of agent smarter than the others, is also supported in Sections (5.1.2) and (5.1.4). Therefore, the forecasting accuracy and portfolio performance are excluded, and what is left is only the saving behavior. Therefore, the significance of risk preference is manifested by the saving behavior.

While earlier studies did recognize the importance of saving behavior to this issue, the attention has been restricted to the rational-equilibrium path. Various

²² One certainly can pursue this further by asking whether we can reverse the result by allowing for a larger degree of difference in forecasting accuracy, and, if so, how much larger a difference is required. We can definitely test for this in a further study, but, for the purpose of negating the absolute importance of forecasting accuracy, this is not necessary.

moments of saving behavior are not a concern as long as they are on the rational-equilibrium path. However, when they are not, this paper shows that moments do have strong implications for survivability. In addition to the first moment, which has already been explicitly noticed by Blume-Easley and Sandroni, we also find the significance of other high-order moments, which are summarized by the box-and-whisker plot. Downside saving rates as well as the dispersion of saving rates can both be important.

Since the saving decision in general is dependent upon the belief, errors in forecasting accuracy can propagate through the saving decision, and manifest themselves in moments of saving. This observation generally applies to all types of agents, except type-1 agents, the log-utility agents, whose saving decisions are independent of their beliefs and are only determined by the exogenously-given discount rate. That is what makes the log-utility agent so different from other agents. The conventional wisdom characterized by the Kelly criterion also works on this, but has not successfully established its validity, in particular, in the general equilibrium context. This paper shows that incorporating learning dynamics is one way of demonstrating its validity.

It is well said that there is only one way to be rational, but an infinite number of ways to be bounded rational. How generally can our findings be extended to other bounded-rational ways, e.g., different learning algorithms, different control parameters, and so on and so forth? This is the type of question frequently asked in this kind of paper. Unfortunately, our knowledge of this is limited. While the idea of error propagation presumably should be shared quite generally by many other learning algorithms, the exact details of implementation may make complete evaluation very difficult, if not impossible.

Nevertheless, we do hope to make a correct claim in relation to our contribution. Our contribution is not to provide a general validity of the relevance of risk preference in a bounded-rational environment, which we have not achieved so far. Rather, our contribution can be considered to be just a point, a starting point maybe, to reflect upon the empirical plausibility of the theoretical result in regard to the irrelevance of risk preference. Agent-based computational modeling is one way of doing this. Would the theoretical result be too strong when applied to the real world? Should risk preference be neglected in further economic analysis of similar issues? Given the findings of this paper, the quest for answers to these topics may have only just begun.

A Appendix

A.1 Evolution at the Low Level: Investment Strategies

A.1.1 Coding and Initialization

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_t^i) and the portfolio (α_t^i) are coded as real-valued numbers:

$$\{\delta_t^i \mid \alpha_{1,t}^i, \alpha_{2,t}^i, \dots, \alpha_{M,t}^i\} \quad (\text{A.1})$$

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

$$GEN_{t,0}^i \equiv \{\delta_{t,n}^i(0), \alpha_{t,n}^i(0)\}_{n=1}^N.$$

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

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

$$(Q_1, Q_2, \dots, Q_M)$$

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

$$\left(\frac{Q_1}{\sum_{q=1}^M Q_q}, \frac{Q_2}{\sum_{q=1}^M Q_q}, \dots, \frac{Q_M}{\sum_{q=1}^M Q_q} \right) \quad (\text{A.2})$$

A.1.2 Fitness Evaluation: Eval $\{ GEN_{t,g}^i \}$

Corresponding to (13), the fitness measure f is simply the H -horizon discounted expected utility:

$$f_t(n, g) \equiv f(\delta_{t,n}^i(g), \alpha_{t,n}^i(g)) \equiv E\left\{ \sum_{h=0}^{H-1} (\beta^i)^h u^i(c_{t+h}^i) \mid B_t^i \right\}, \quad (\text{A.3})$$

where $f_t(n, g)$ refers to the fitness of the n th investment strategy in the population $GEN_{t,g}^i$ (i.e. the g th generation of the GA cycle). The Monte Carlo simulation technique is used to evaluate the fitness (A.3). The way to do so is to simulate a certain 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 (A.3), i.e.

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

Then, we estimate $f_t(n, g)$ by taking the sample average,

$$\hat{f}_t(n, g) = \frac{\sum_{l=1}^L \sum_{h=0}^{H-1} (\beta^i)^h U^i(c_{t+h}^i | l)}{L}. \quad (\text{A.4})$$

A.1.3 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 (A.4).

$$\text{Eval} : \{ \delta_{t,n}^i(g), \alpha_{t,n}^i(g) \}_{n=1}^N \rightarrow \{ f_t(n, g) \}_{n=1}^N \quad (\text{A.5})$$

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 them, 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 p_{cross} (*crossover rate*), 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.

$$\begin{aligned} I_z &\equiv (\delta_{t,z}^i(g), \alpha_{t,z}^i(g)) \\ &= \frac{f_t(x, g)}{f_t(x, g) + f_t(y, g)} (\delta_{t,x}^i(g), \alpha_{t,x}^i(g)) + \frac{f_t(y, g)}{f_t(x, g) + f_t(y, g)} (\delta_{t,y}^i(g), \alpha_{t,y}^i(g)) \end{aligned} \quad (\text{A.6})$$

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

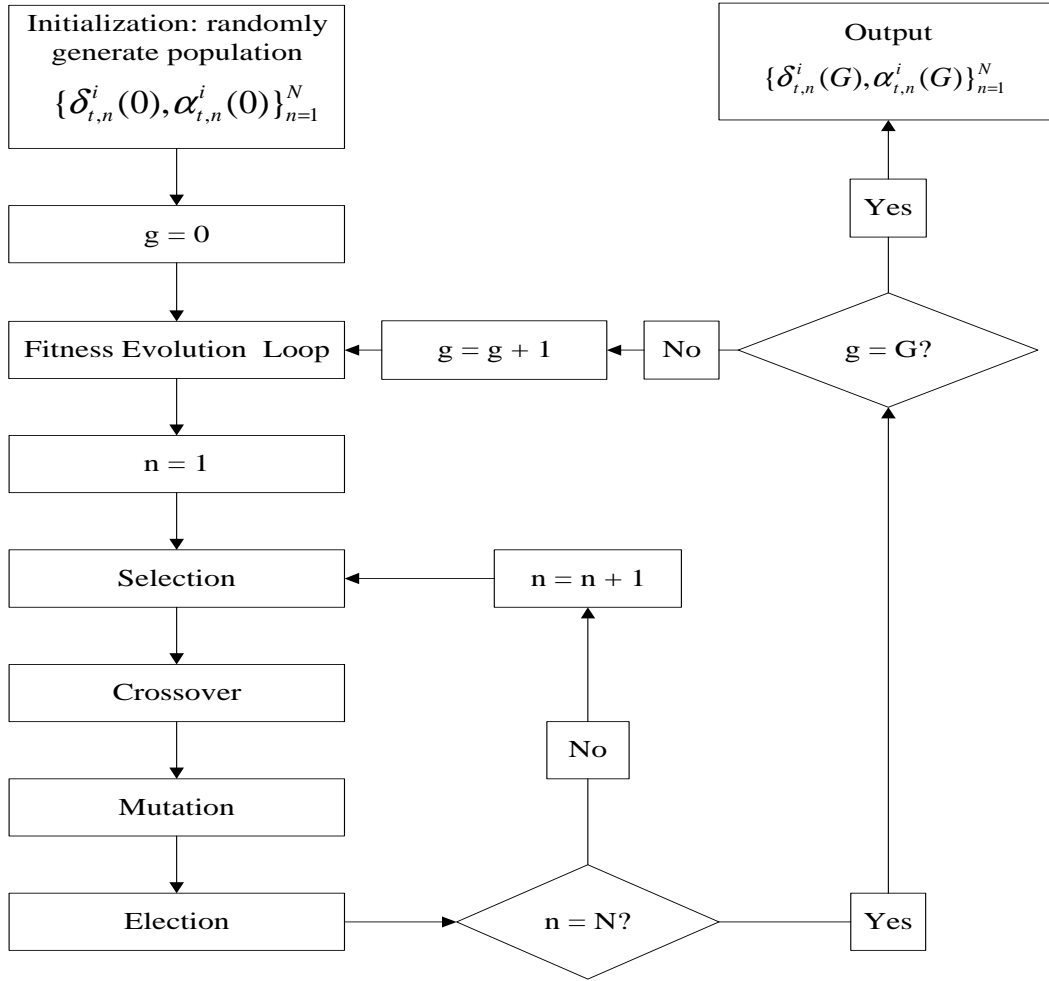


Fig. A.1. Flowchart of the Low-Level GA

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,

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

randomly generated from $U(0, 1)$, will replace $\alpha_{t,z}^i(g)$. Then the rescaling technique described in (A.2) will be applied. We call the resultant strategy $I_{z'}$.

Election: The use of the election operator examines whether the new investment strategy is expected to perform better than the one it replaced. In election, we shall use (A.4) 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, $GEN_{t,g+1}^i$.

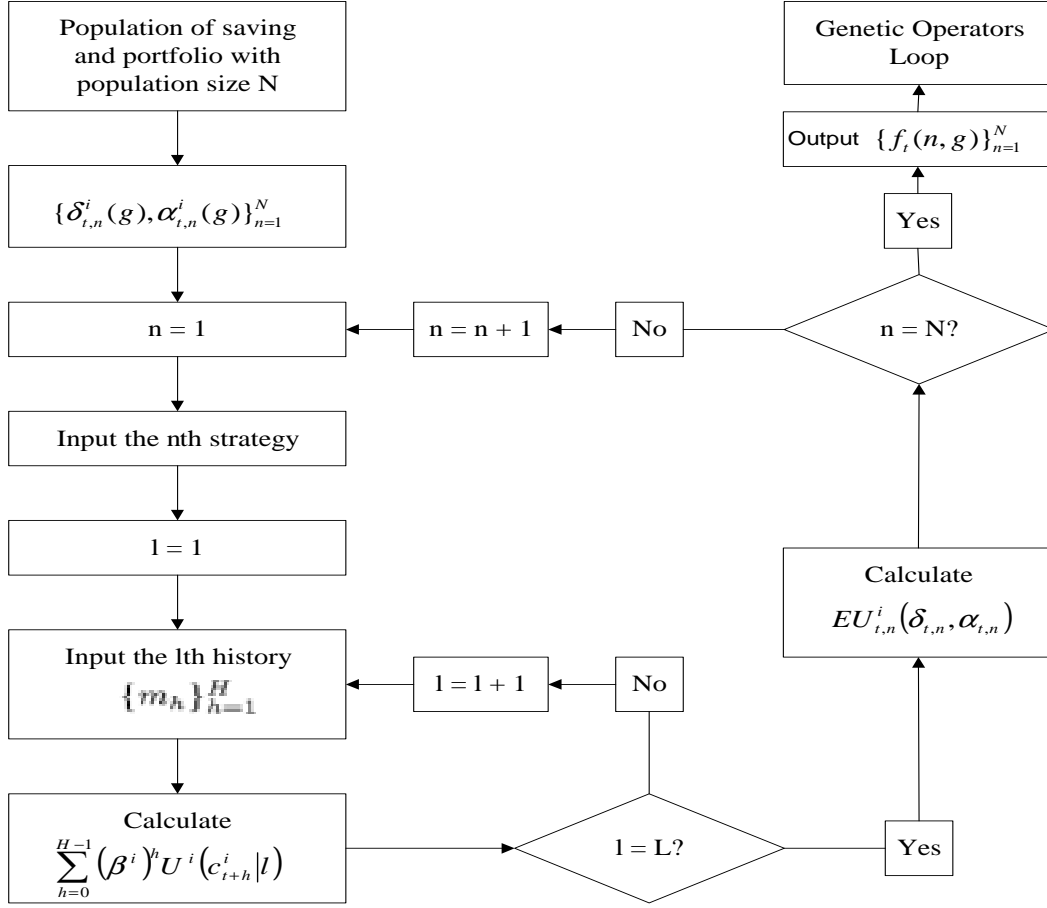


Fig. A.2. The Flowchart of the Investment Optimization

A.1.4 Loops

Once a new investment strategy is generated, a *loop* (Figure A.1) 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 $GEN_{t,g+1}^i$ are generated. $GEN_{t,g+1}^i$ will be evaluated based on the **Eval** procedure, and based on the evaluation, genetic operators will be applied to $GEN_{t,g+1}^i$ to generate $GEN_{t,g+2}^i$. 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, $GEN_{t,G}^i$.

$$(\delta_t^{i,*}, \alpha_t^{i,*}) = \arg \max_{GEN_{t,G}^i} \{f_t(n, G)\}_{n=1}^N \quad (\text{A.7})$$

A.2 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* $\{B_{j,t}^i\}_{j=1}^J$. 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* ([12], [2], [23], [8], [1]). 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.

A.2.1 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 k th-order Markov process. According to this belief, he would use a part of the historical data $\{m_{t-s}\}_{s=v+1}^{v+d+1}$, referred as to the *training period*, to estimate the Markov transition matrix, and the rest of the data $\{m_{t-s}\}_{s=1}^v$, referred to as the *validation period*, to validate the estimated model. As a result, each belief can be represented by a binary string, of length $\tau_1 + \tau_2$,

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

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

$$k = \left(\sum_{i=1}^{\tau_1} 2^{\tau_1-i} a_i \right) \quad (\text{A.8})$$

over the last

$$d = \left(\sum_{i=\tau_1+1}^{\tau_1+\tau_2} 2^{\tau_1+\tau_2-i} a_i \right) + c \quad (\text{A.9})$$

days. To facilitate estimation, d cannot be too small, and that demands an additional constant of c . 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 first-order Markov.

At the initial date ($t = 0$), all investors are endowed with a population of J

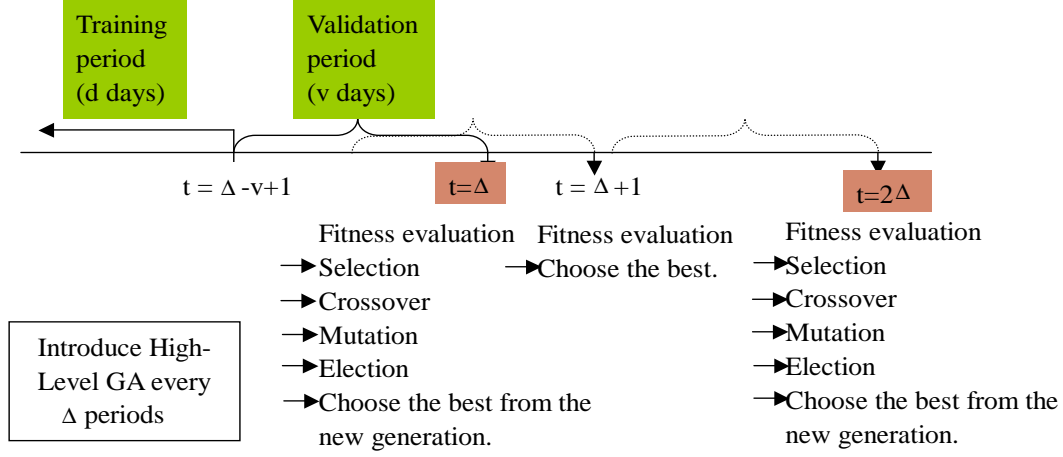


Fig. A.3. The Belief Updating Scheme

beliefs, which are randomly generated. Then every Δ days, this population of belief will be reviewed and revised based on the fitness function, which is a kind of likelihood function to be specified below.

A.2.2 Belief Updating Scheme

Agents in our model follow the practice of machine learning. They are supposed to care about the risk of over-fitting, and hence use data in the validation period to perform 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. The belief updating scheme is outlined in Figure A.3.

The essence of the belief updating scheme is to maintain a style of on-line learning, while not to overload the computational intensity. 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 a belief $B_{j,t}^i$ is its associated *likelihood*, evaluated by the validation set $\{m_{t-s}\}_{s=1}^v$,

$$L_{j,t}^i = L(\{m_{t-s}\}_{s=1}^v | B_{j,t}^i), \quad (\text{A.10})$$

Equation (A.10) is the likelihood of the observations $\{m_{t-s}\}_{s=1}^v$ in the validation period under the belief $B_{j,t}^i$. Every Δ periods, after they finish the evaluation of each belief's fitness, they apply the genetic operation to update their belief set (see Section A.2.3), and the belief with the highest fitness will be chosen. Even in the period that the genetic operation is not applied, say when $t \in [\Delta + 1, 2\Delta - 1]$, they evaluate the fitness of beliefs in their current belief set using the newest data and choose the best from it.

A.2.3 Genetic Operation

Once the procedure of evaluating each belief's fitness (**Eval** $\{B_{j,t-1}^i\}_{j=1}^J$) is completed, all beliefs are associated with a fitness which is the output of (A.10).

$$\mathbf{Eval} : \{B_{j,t-1}^i\}_{j=1}^J \rightarrow \{L_{j,t-1}^i\}_{j=1}^J \quad (\text{A.11})$$

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 p_{cross} , 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 $B_{x,t-1}^i = 0010101010$, corresponding to a belief of $(k_x, d_x) = (0, 170)$, and $B_{y,t-1}^i = 0111110010$, corresponding to $(k_y, d_y) = (0, 498)$. Then, an offspring, B_z^i , can be

$$B_z^i = 0011100010 \rightarrow (k_z, d_z) = (0, 226).$$

Mutation: There is a small probability p_{mutate} (mutation rate) by which each bit of B_z^i may encounter a change. For example, the mutation which changes the fifth bit from "1" to "0", and the last bit from "0" to "1" will result in a new string:

$$B_{z'}^i = 0011000011 \rightarrow (k_{z'}, d_{z'}) = (0, 195).$$

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

A.2.4 Loops

Once a belief is generated, a loop in Figure A.4 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 $\{B_{j,t}^i\}_{j=1}^J$ are generated. One of the beliefs, $B_{j,t}^{i,*}$, will be chosen based on the likelihood criteria,

$$B_t^{i,*} = \arg \max_j L(\{m_{t-s}\}_{s=1}^v | B_{j,t}^i), \quad (\text{A.12})$$

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

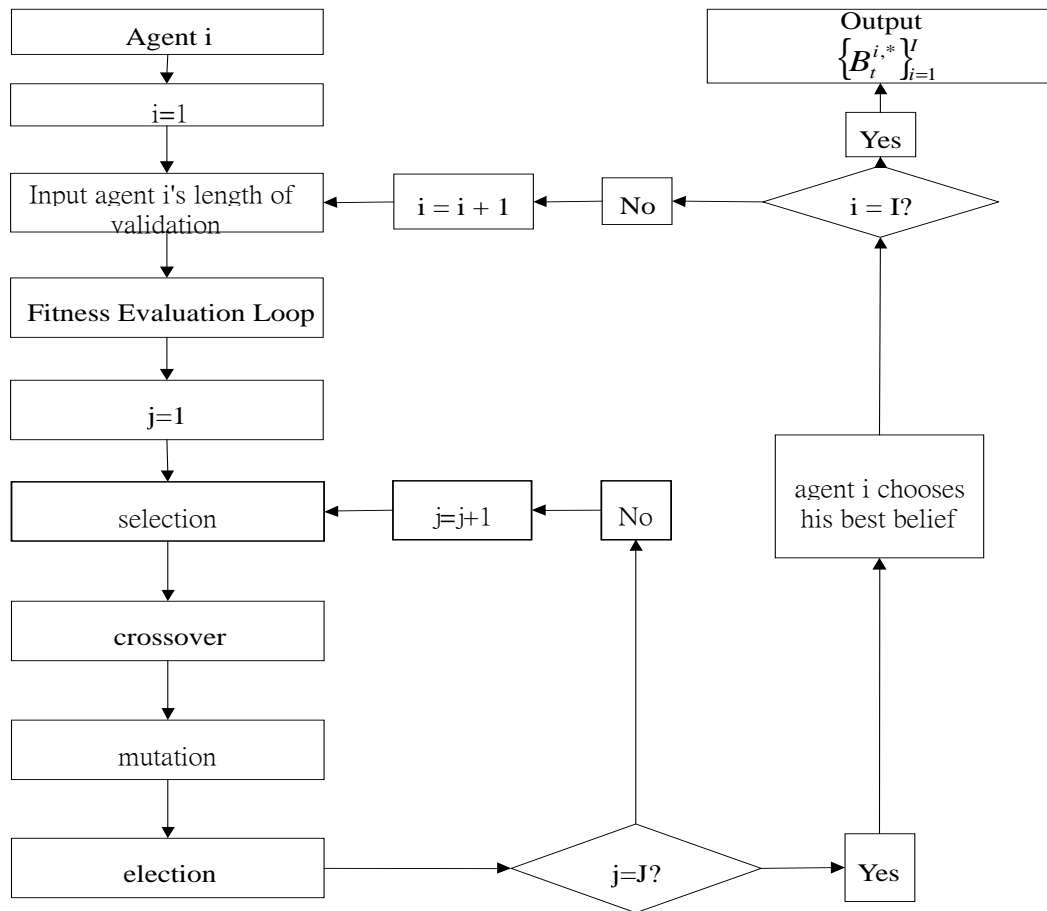


Fig. A.4. Flowchart of the High-Level GA

References

- [1] Arifovic, J. and Maschek, M., 2003. Expectations and Currency Crisis—An Experimental Approach. Paper presented at the *9th International Conference on Computing in Economics and Finance*, University of Washington, Seattle, July 11-13.
- [2] Arthur, W.B., Holland, J. , LeBaron, B. , Palmer, R. and Tayler, P., 1997. Asset Pricing under Endogenous Expectations in an Artificial Stock Market, In: W.B. Arthur, S. Durlauf and D. Lane, (Eds.), *The Economy as an Evolving Complex System II*, Addison-Wesley, Reading, MA, pp. 15-44.
- [3] Barberis, N. and Thaler, R., 2002. A Survey of Behavioral Finance, in: G. Constantinides, M. Harris and R. Stulz (Eds.), *Handbook of the Economics of Finance*, pp. 1051-1121.
- [4] Blume, L., Easley, D., 1992. Evolution and Market Behavior. *Journal of Economic Theory* 58, 9-40.
- [5] Blume, L., Easley, D., 2001. If You're So Smart, Why Aren't You Rich? Belief Selection in Complete and Incomplete Markets. Working paper.
- [6] Bullard, J., Duffy, J., 1999. Using Genetic Algorithms to Model the Evolution of Heterogenous Beliefs. *Computational Economics* 13, 41-60.

- [7] Chen, S.-H., Huang, Y.-C., 2004. Risk Preference and Survival Dynamics. Working Paper Series 2004-2, AI-ECON Research Center, National Chengchi University.
- [8] 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.
- [9] Feldman, J., 1962. Computer Simulation of Cognitive Processes, in: H. Boroko (Ed.), *Computer Applications in the Behavioral Sciences*, Prentice Hall.
- [10] Friend, I., Blume, M.E., 1975. The Demand for Risky Assets. *American Economic Review* 65, 900-922.
- [11] Gordon, M.J., Paradis, G.E., Rorke, C.H., 1972. Experimental Evidence on Alternative Portfolio Decision Rules. *American Economic Review* 62, 107-118.
- [12] Holland, J., Miller, J., 1991. Artificial Adaptive Agents in Economic Theory. *American Economic Review* 81, 365-370.
- [13] Huang, C.F., Litzenberger, R.H., 1988. *Foundations for Financial Economics* (Prentice Hall)
- [14] Kelly, J.L., 1956. A New Interpretation of Information Rate. *Bell System Technical Journal* 35, 917-926.
- [15] Laibson, D., 1998. Life-Cycle Consumption and Hyperbolic Discount Functions. *European Economic Review* 42, 861-871.
- [16] Lettau, M., 1997. Explaining the Facts with Adaptive Agents: the Case of Mutual Fund Flows. *Journal of Economic Dynamics and Control* 21, 1117-1147.
- [17] Rabin, M., 1998. Psychology and Economics. *Journal of Economic Literature* 36, 11-46.
- [18] Sandroni, A., 2000. Do Markets Favor Agents Able to Make Accurate Predictions? *Econometrica* 68, 1303-1341.
- [19] Sciubba, E., 1999. The Evolution of Portfolio Rules and the Capital Asset Pricing Model. DAE Working Paper n. 9909, University of Cambridge.
- [20] Spear, S., 1989. Learning Rational Expectations under Computability Constraints. *Econometrica* 57, 889-910.
- [21] Tay, N., Linn, S., 2001. Fuzzy Inductive Reasoning, Expectation Formation and the Behavior of Security Prices. *Journal of Economic Dynamics and Control* 25, 321-361.
- [22] Tesfatsion, L., 2001. Introduction to the Special Issue on Agent-Based Computational Economics. *Journal of Economic Dynamics and Control* 25, 281-293.
- [23] 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* 24, 1-19.