An Agent Market Model Using Evolutionary Game Theory

Benoît S. Montin*[†]   Craig A. Nolder[†]

Department of Mathematics - Florida State University

JEL classification codes: C73, G11.

Key words: artificial market, evolutionary game theory, stochastic equilibrium.

Abstract

Stock price fluctuations result from interactions between economic agents. Modelling the financial world as a complex self-organizing system is thus natural. To keep tractability, obtain theoretical results and develop our intuition about complexity economics, we have constructed a very simple artificial stock market. Despite its necessary over-simplification, our model is rich enough to offer interesting results about limit behavior and suggests monetary policies. Our multi-agent model also exhibits real world features that more traditional financial models usually fail to explain or consider.

*Address: 208 Love Building, Tallahassee FL 32306-4510. Phone: (850) 644-2202. Fax: (850) 644-4053. Email: bmontin@math.fsu.edu

[†]The authors are grateful to Dr. Mike Mesterton-Gibbons and Dr. Paul M. Beaumont for their assistance.

1 Introduction

First, it is natural to ask the following question. Why shall we try to develop a new methodology? Indeed, there already exist skillful models where agents maximize their expected utility by forming rational expectations about future outcomes based on past observations of the economy in equilibrium[1]. But what if the market is new or perturbed by external factors? Macroeconomic and financial external factors are shocks that might put the economy out of equilibrium. Evolutionary models allow studying how a composite system reacts to such disturbances. For example, it becomes possible to study the dynamics of the economy in response to a sequence of monetary policies. In this context, evolutionary game theory can strongly contribute to our economic understanding.

We have constructed a simplistic model. Three risk-averse economic agents interact in our artificial market. An agent can be any financial institution or a single investor. Time is discrete. There are two kinds of assets: a stock and zero coupon bonds earning a riskless rate of interest r. Bonds are issued at each time step and are of maturity one period. The role of the bonds is similar to the one of cash paying interest rates in saving accounts. For convenience, we adopt the following convention: all bonds have the same time zero value denoted by B_0. Buying a bond at time t_k thus represents a cost at time t_k of $B_k = B_0(1+r)^{t_k}$ and a payoff at time t_{k+1} of $B_{k+1} = B_0(1+r)^{t_{k+1}}$. At each time step, an agent can own bonds only or stocks only. Agents simply seek the maximization of their wealth: there is no consumption.

[1]new classical school

The state of the composite system, at a given time, can be represented as in figure 1. The $x-$axis corresponds to the proportion of wealth that an agent invests in the stock. The wealth is given by the $y-$component. With our restrictions, an agent can only be in one of two states:

- state $|0\rangle$: agent n^o i owns bonds only ($x = 0$, $y \geq 0$),

- state $|1\rangle$: agent n^o i owns stocks only ($x = 1$, $y \geq 0$).

The quantum mechanical formalism is only used as a convenient mathematical tool. In quantum game theory, players can adopt quantum strategies usually more efficient than classical mixed strategies (e.g., Lee and Johnson, 2003). We do not consider quantum strategies in this article.

Figure 1: the economic space

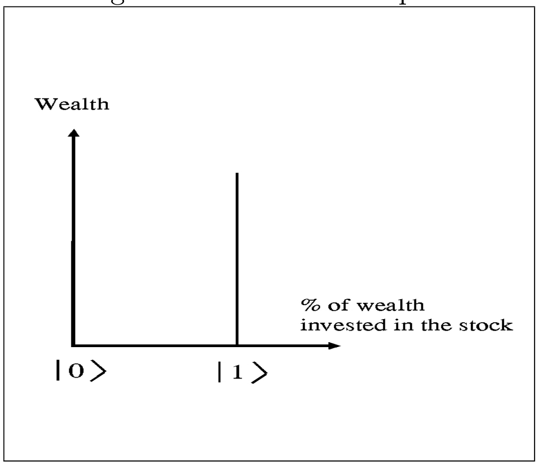

There are six configurations of interest distinguishing all possible forms of allocations (figure 2). For example, in configuration A, agents n^o 1 and n^o 2 own bonds, agent n^o 3 owns stocks. The wealths could be differently configured. Stock price fluctuations result from transitions from one configuration to another. We do not consider the two static configurations where all agents own bonds only or stocks only. Indeed, for all agents to be in the state $|0\rangle$ (respectively $|1\rangle$), it must be so all the time.

Figure 2: the six configurations

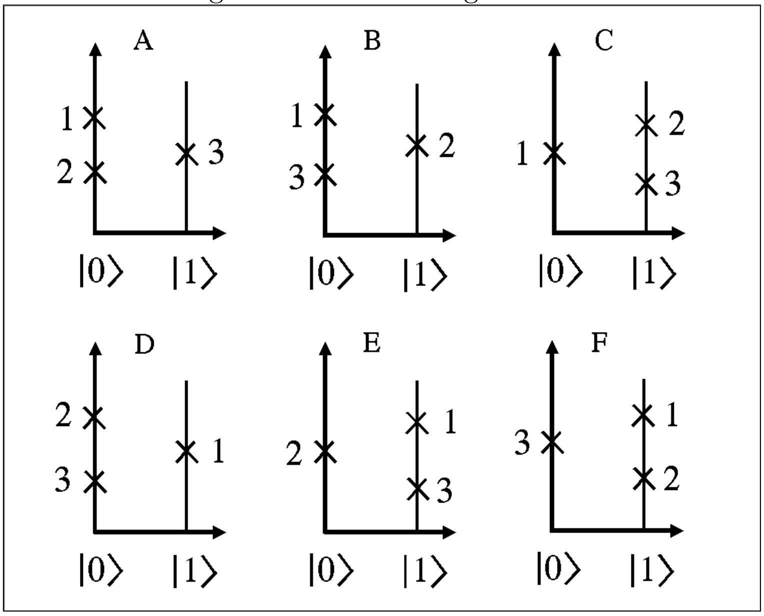

It is important to notice that all transitions are not feasible. If our simplified economy is in configuration C at time t_k, then configuration A cannot be attained at time t_{k+1}. Indeed, there is no potential buyer of stock shares to allow agent n^o 2 moving from the state $|1\rangle$ to the state $|0\rangle$. Feasible transitions are schematized in figure 3: if the system is in configuration A at time t_k, then it can remain in configuration A or switch to the configurations B, D, or F at time t_{k+1}.

When a transition occurs, the budget constraints and conservation laws allow calculating the new stock price and allocations. Assume for example that our simplified economy is in configuration A at time t_k and in configu-

Figure 3: feasible transitions

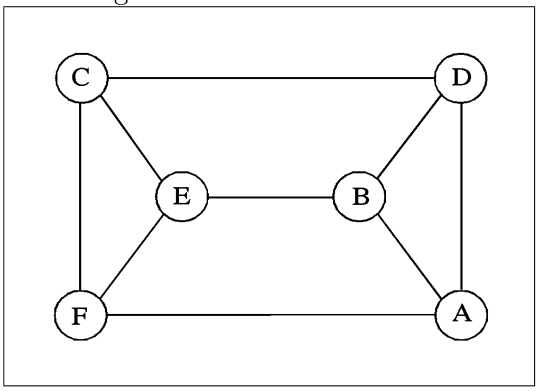

ration F at time t_{k+1}. We use the following notations:

$$\begin{cases} S_k \text{ denotes the stock price at time } t_k \\ a_k^{(i)} \text{ is the number of stock shares owned by agent } n^o\ i \text{ at time } t_k \\ b_k^{(i)} \text{ is the number of bond shares owned by agent } n^o\ i \text{ at time } t_k \\ W_k^{(i)} \text{ represents agent } n^o\ i\text{'s wealth at time } t_k \end{cases}$$

If the system is in configuration A at time t_k, then the respective wealths are:

$$\begin{cases} W_k^{(1)} = b_k^{(1)} B_0 (1+r)^{t_k} \\ W_k^{(2)} = b_k^{(2)} B_0 (1+r)^{t_k} \\ W_k^{(3)} = a_k^{(3)} S_k \end{cases} \tag{1}$$

The budget constraints and conservation law are:

- Our economy being closed, the total number of stock shares remain constant:

$$a_{k+1}^{(1)} + a_{k+1}^{(2)} = a_k^{(3)} \tag{2}$$

- Agents n^o 1 and n^o 2 buy stocks from the earnings of selling bonds:

$$b_k^{(1)} B_0 (1+r)^{t_{k+1}} = a_{k+1}^{(1)} S_{k+1} \tag{3}$$

$$b_k^{(2)} B_0 (1+r)^{t_{k+1}} = a_{k+1}^{(2)} S_{k+1} \tag{4}$$

- Agent n^o 3 buys bonds from the earnings of selling stocks:

$$a_k^{(3)} S_{k+1} = b_{k+1}^{(3)} B_0 (1+r)^{t_{k+1}} \tag{5}$$

The stock price at time t_{k+1} is thus:

$$S_{k+1} = \frac{b_k^{(1)} + b_k^{(2)}}{a_k^{(3)}} B_0 (1+r)^{t_{k+1}} \tag{6}$$

and the new wealths are:

$$\begin{cases} W_{k+1}^{(1)} = \left(\frac{b_k^{(1)}}{b_k^{(1)}+b_k^{(2)}} a_k^{(3)}\right) S_{k+1} \\ W_{k+1}^{(2)} = \left(\frac{b_k^{(2)}}{b_k^{(1)}+b_k^{(2)}} a_k^{(3)}\right) S_{k+1} \\ W_{k+1}^{(3)} = (b_k^{(1)} + b_k^{(2)}) B_0 (1+r)^{t_{k+1}} \end{cases} \tag{7}$$

Similar calculations can be carried for all feasible transitions. To simplify, we can form two classes of configurations. For the first class, two agents are in the state $|0\rangle$ and one agent is in the state $|1\rangle$. The first class thus regroups the configurations A, B and D. The remaining configurations belong to the second class. It is sufficient to consider transitions occurring from one configuration of each class to the respective attainable configurations. The other cases are obtained by permutations of the agents.

2 How do transitions occur?

As in game theory, agents adopt mixed strategies constructed from the two pure strategies:

- hold bonds if already own them, otherwise try to trade,

- hold stocks if already own them, otherwise try to trade.

Strategies can be modelled with two special orthogonal matrices R (Remain) and C (Change) defined as:

$$R = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad C = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \tag{8}$$

Recall that a square matrix U of dimension n is a unitary matrix if:

$$U^\dagger U = I_n \tag{9}$$

$$\text{where} \begin{cases} U^\dagger \text{ represents the conjugate transpose of } U \\ I_n \text{ is the n-dimensional identity matrix} \end{cases}$$

The above unitary operators transform collapsed wave functions as suggested by their names:

$$C\,|0\rangle = C \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ -1 \end{pmatrix} = -|1\rangle \tag{10}$$

$$C\,|1\rangle = C \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix} = |0\rangle \tag{11}$$

And trivially, $R\,|0\rangle = |0\rangle$ and $R|1\rangle = |1\rangle$.

Let $p_{0,k}^{(i)}$ (respectively $p_{1,k}^{(i)}$) denote agent n^o i's probability to play the first (respectively second) pure strategy at time t_k. Naturally, $p_{0,k}^{(i)} + p_{1,k}^{(i)} = 1$. Assume that agent n^o i is in the state $|0\rangle$ at time t_k and adopts the mixed strategy:

$$U_k^{(i)} = \sqrt{p_{0,k}^{(i)}}\,R + \sqrt{p_{1,k}^{(i)}}\,C \tag{12}$$

It is straightforward to check that $U_k^{(i)}$ is a unitary operator. Agent n^o i's wave function then becomes at time t_{k+1}:

$$\begin{aligned} |\psi_{k+1}^{(i)}\rangle &= (\sqrt{p_{0,k}^{(i)}}\,R + \sqrt{p_{1,k}^{(i)}}\,C)|0\rangle \\ &= \sqrt{p_{0,k}^{(i)}}\,|0\rangle - \sqrt{p_{1,k}^{(i)}}\,|1\rangle \end{aligned} \tag{13}$$

In other words, agent n^o i is in the state $|0\rangle$ (respectively in the state $|1\rangle$) with a probability $p_{0,k}^{(i)}$ (respectively $p_{1,k}^{(i)}$).

The state of the composite system at time t_{k+1} is given by the Kronecker tensor product $\otimes$:

$$|\psi_{k+1}^{(1)}\rangle \otimes |\psi_{k+1}^{(2)}\rangle \otimes |\psi_{k+1}^{(3)}\rangle = (U_k^{(1)}|\psi_k^{(1)}\rangle) \otimes (U_k^{(2)}|\psi_k^{(2)}\rangle) \otimes (U_k^{(3)}|\psi_k^{(3)}\rangle) \quad (14)$$

where $|\psi_k^{(i)}\rangle$, for $i \in \{1,2,3\}$, are collapsed wave functions ($|\psi_k^{(i)}\rangle = |0\rangle$ or $|\psi_k^{(i)}\rangle = |1\rangle$). The Kronecker tensor product $\otimes$ of two vectors $\mathbf{x}$ and $\mathbf{y}$ is the larger vector formed from all possible products of the elements of $\mathbf{x}$ with those of $\mathbf{y}$. The elements are arranged in the following order:

$$\begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} \otimes \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_p \end{pmatrix} = \begin{pmatrix} x_1 y_1 \\ x_1 y_2 \\ \vdots \\ x_1 y_p \\ \vdots \\ x_n y_p \end{pmatrix} \quad (15)$$

At time t_{k+1}, offers for potential trades or refusals are made by the agents. If a trade occurs, portfolio allocations are modified. Otherwise, they remain the same as at time t_k. In both cases, portfolio allocations are observed at time t_{k+1}. In other words, the individual wave functions $|\psi_{k+1}^{(i)}\rangle$ collapse to one of the two pure states $|0\rangle$ or $|1\rangle$. The components of $|\psi_{k+1}^{(1)}\rangle \otimes |\psi_{k+1}^{(2)}\rangle \otimes |\psi_{k+1}^{(3)}\rangle$ give the probabilities of occurrence of each configuration at time t_{k+1}. The probabilities associated to non-feasible transitions are naturally added to the probability that the system remains in its present configuration.

To complete the description of our artificial stock market, we now explain how agents choose their mixed strategies $p_{0,k}^{(i)}$. Replicator dynamics provide the basis for realistic behavioral rules (Canning, 1992):

> "A more intuitively appealing approach is to assume that agents follow behavioral rules that are less complex than full Bayesian learning. A behavioral rule should tell the agent what actions to take, and how to change these actions in the light of experience."

Replicator dynamics originated in the field of evolutionary biology. They have recently become very popular among evolutionary game theorists (e.g., Fudenberg and Levine, 1998, Samuelson, 1997). In their discrete version, the replicator equations can be expressed as:

$$p_{0,k+1}^{(i)} = p_{0,k}^{(i)} \left\{ 1 + \frac{\gamma\,\Delta t\,[u_k^{(i)} - \overline{u}_k^{(i)}]}{1 + \gamma\,\Delta t\,\overline{u}_k^{(i)}} \right\} \tag{16}$$

where

$$\begin{cases} \gamma \text{ is a learning rate} \\ \Delta t = t_{k+1} - t_k \text{ represents the time interval between two potential trades} \\ u_k^{(i)} \text{ is the payoff to using the first pure strategy} \\ \overline{u}_k^{(i)} \text{ is the average expected payoff} \end{cases}$$

Learning is myopic and local: the probability of adopting the first pure strategy is directly related to how well the strategy has been doing in the past.

Agent n^o i's payoff at time t_{k+1} is defined as the discounted increase of wealth:

$$P_{k+1}^{(i)} = \frac{W_{k+1}^{(i)}}{(1+r)^{t_{k+1}}} - \frac{W_k^{(i)}}{(1+r)^{t_k}} \tag{17}$$

We can easily calculate the payoffs associated to each feasible transition.

For example, if the system is in configuration A at time t_k, then the payoff matrices are:

- If agent n^o 3 adopts the first pure strategy:

		agent n^o 2	
		1^{st}	2^{nd}
agent n^o 1	1^{st}	$\left(0,0,-a_k^{(3)}S_k^*\left(1-\frac{1}{(1+r)^{t_{k+1}-t_k}}\right)\right)$	$(0,0,b_k^{(2)}B_0-a_k^{(3)}S_k^*)$
	2^{nd}	$(0,0,b_k^{(1)}B_0-a_k^{(3)}S_k^*)$	$(0,0,(b_k^{(1)}+b_k^{(2)})B_0-a_k^{(3)}S_k^*)$

- If agent n^o 3 adopts the second pure strategy:

		agent n^o 2	
		1^{st}	2^{nd}
agent n^o 1	1^{st}	$\left(0,0,-a_k^{(3)}S_k^*\left(1-\frac{1}{(1+r)^{t_{k+1}-t_k}}\right)\right)$	$\left(0,0,-a_k^{(3)}S_k^*\left(1-\frac{1}{(1+r)^{t_{k+1}-t_k}}\right)\right)$
	2^{nd}	$\left(0,0,-a_k^{(3)}S_k^*\left(1-\frac{1}{(1+r)^{t_{k+1}-t_k}}\right)\right)$	$\left(0,0,-a_k^{(3)}S_k^*\left(1-\frac{1}{(1+r)^{t_{k+1}-t_k}}\right)\right)$

We notice that the payoff to buying stocks is zero. However, it is easy to check that owning stocks for several periods can be a winning strategy. Suppose for example that the composite system is in configuration A at time t_k. If all agents participate in a trade at time t_{k+1}, then the system switches to configuration F and the stock price and allocations become:

$$\begin{cases} S_{k+1}=\frac{b_k^{(1)}+b_k^{(2)}}{a_k^{(3)}}B_0(1+r)^{t_{k+1}} \\ b_{k+1}^{(1)}=b_{k+1}^{(2)}=0 \quad b_{k+1}^{(3)}=b_k^{(1)}+b_k^{(2)} \\ a_{k+1}^{(1)}=\frac{b_k^{(1)}}{b_k^{(1)}+b_k^{(2)}}a_k^{(3)} \quad a_{k+1}^{(2)}=\frac{b_k^{(2)}}{b_k^{(1)}+b_k^{(2)}}a_k^{(3)} \quad a_{k+1}^{(3)}=0 \end{cases} \tag{18}$$

If agent n^o 1 decides to keep the current portfolio (stocks only) for a second period: $a_{k+2}^{(1)}=a_{k+1}^{(1)}$, and if agent n^o 2 and n^o 3 participate in a trade at time t_{k+2}, then the stock price S_{k+2} is:

$$S_{k+2} = \frac{b_{k+1}^{(3)}}{a_{k+1}^{(2)}} B_0 (1+r)^{t_{k+2}} = \frac{(b_k^{(1)} + b_k^{(2)})^2}{b_k^{(2)} a_k^{(3)}} B_0 (1+r)^{t_{k+2}} \tag{19}$$

Agent n^o 1's wealth at time t_{k+2} is thus:

$$W_{k+2}^{(1)} = a_{k+2}^{(1)} S_{k+2} = \frac{b_k^{(1)}}{b_k^{(2)}} (b_k^{(1)} + b_k^{(2)}) B_0 (1+r)^{t_{k+2}} \tag{20}$$

Focusing on these two periods only and assuming the above sequence of events, having bought stocks at time t_{k+1} and kept them at time t_{k+2} makes a better strategy for agent n^o 1 than always owning bonds. Such a strategy indeed induces greater wealth:

$$\frac{b_k^{(1)}}{b_k^{(2)}} (b_k^{(1)} + b_k^{(2)}) B_0 (1+r)^{t_{k+2}} \geq b_k^{(1)} B_0 (1+r)^{t_{k+2}} \tag{21}$$

From the replicator dynamics, an agent who experiences such a winning strategy is more likely to own stocks in the future. As we will discuss later, there are other reasons why stocks might be attractive.

Finally, to make sure that the discrete dynamic equations indeed provide probabilities, that is to say:

$$\forall k \in \mathbf{N} : \; 0 \leq p_{0,k}^{(i)} \leq 1 \tag{22}$$

we impose on $\gamma \, \Delta t$ the constraint:

$$\gamma \, \Delta t \, (\sum_{i=1}^{3} a_0^{(i)}) \, M < 1 \tag{23}$$

where M is an upper-bound for the discounted stock price. If t_x denotes the first time at which two agents own stocks, then M can be defined as (Montin, 2004):

$$M = \max \left\{ S_0, \frac{\sum_{i=1}^{3} b_0^{(i)}}{\min_{\{i | a_x^{(i)} \neq 0\}} \{a_x^{(i)}\}} B_0 \right\} \tag{24}$$

3 Stochastic equilibria

> "For dynamic economic models, an equilibrium (or steady state) is defined to be a point in the state space that is stationary under the period to period transition rule. In the case of stochastic economies, a state cannot be stationary in the same sense as that of deterministic models, given that shocks continue to disturb activity in each period. Instead a steady state must be viewed as a situation where the probabilistic laws that govern the state variables cease to change over time (Stachurski)."

The state s_k of the composite system at time t_k is completely described by the following random vector:

$$s_k = (x_k, p_{0,k}^{(1)}, p_{0,k}^{(2)}, p_{0,k}^{(3)}, a_k^{(1)}, a_k^{(2)}, a_k^{(3)}, b_k^{(1)}, b_k^{(2)}, b_k^{(3)}, S_k^*) \tag{25}$$

where x_k equals A, B, C, D, E or F depending on the present configuration at time t_k and S_k^* denotes the discounted stock price at time t_k. Let's call Σ the state space. By construction:

$$\Sigma \subset \{a, \cdots, F\} \times [0,1]^3 \times [0, \sum_{i=1}^{3} a_0^{(i)}]^3 \times [0, \sum_{i=1}^{3} b_0^{(i)}]^3 \times [0, M] \tag{26}$$

We consider the square metric ρ on the state space Σ:

$$\rho(\mathbf{u}, \mathbf{v}) = \max_i \{d_i(u_i, v_i)\} \tag{27}$$

The metric chosen for the first dimension induces the discrete topology:

$$\begin{aligned} d_1(u_1, v_1) &= 1 \text{ if } u_1 \neq v_1 \\ &= 0 \text{ if } u_1 = v_1 \end{aligned} \tag{28}$$

The usual distance on $\mathbf{R}$ is used for the other dimensions. We complete the state space Σ with its limit points so that it is compact. Let $\mathbf{\Sigma}$ denote the

Borel $\sigma-$field of Σ.

Agents face the following set of events:

$$\begin{aligned} e_1 &= \text{configuration } A \text{ is present at the following time step} \\ &\vdots \\ e_6 &= \text{configuration } F \text{ is present at the following time step} \end{aligned}$$

Let $E = \{e_1, \cdots, e_6\}$ be the space of events and $\mathbf{E}$ be its power set.
The stochastic kernel $Q:\ \Sigma \times \mathbf{E} \to [0,1]$ gives the probability $Q(s, A)$ of realizing the event $A \in \mathbf{E}$ given that the current state is $s \in \Sigma$.
The mapping $\theta:\ \Sigma \times E \to \Sigma$ specifies which state succeeds: $s_{k+1} = \theta(s_k, e_i)$.
We could write explicitly Q and θ (Montin, 2004). With the above definitions, it is easy to prove that our simplified economy is a random dynamical system (Futia, 1982). As such, the function $P:\ \Sigma \times \mathbf{\Sigma} \to [0,1]$ defined by:

$$P(s, A) = Q(s, (\theta^{-1}A)_s) \tag{29}$$

where $(\theta^{-1}A)_s = \{e \in E \,|\, \theta(s,e) \in A\}$ is the section of $\mathrm{Im}\theta \cap A$ determined by s, is a transition probability. With our choice of state representation, our dynamic economy is thus a discrete-time Markov process.

The $n-$step transition probability $P^n(s, A)$ is defined recursively:

$$P^0(s, A) = \begin{cases} 1 & \text{if } s \in A \\ 0 & \text{otherwise} \end{cases} \qquad P^n(s, A) = \int_\Sigma P^{n-1}(s, dt)\, P(t, A) \tag{30}$$

Let $B(\Sigma)$ denote the set of all bounded, $\mathbf{\Sigma}-$measurable, real-valued functions on Σ. $B(\Sigma)$ is a Banach space under the sup norm: $\|f\| = \sup_{s\in\Sigma} |f(s)|$.
The Markov operator associated to the transition probability P is the continuous linear transformation $T:\ B(\Sigma) \to B(\Sigma)$ defined by:

$$(T\, f)(s) = \int_\Sigma f(t) P(s, dt) \tag{31}$$

Notice that: $\|T\| = \sup_{\|f\|\leq 1} \|T\,f\| = 1$.

By isometric isomorphism (Dunford and Schwartz, 1957), we can define the adjoint

$T^* : ba(\Sigma) \to ba(\Sigma)$ by:

$$\forall A \in \mathbf{\Sigma}, \quad (T^*\lambda)(A) = \int_\Sigma P(s,A)\lambda(ds) \tag{32}$$

where $ba(\Sigma)$ denotes the Banach space of all bounded finitely additive set functions λ under the total variation norm. Notice that T^* maps probabilities into probabilities.

More generally, the n^{th} iterate satisfies: $[(T^*)^n\lambda](A) = \int_\Sigma P^n(s,A)\lambda(ds)$.

By definition, an invariant probability measure λ, for a transition probability P, satisfies:

$$\forall A \in \mathbf{\Sigma} : \quad \int_\Sigma P(s,A)\lambda(ds) = \lambda(A) \tag{33}$$

In other words, λ is a fixed point of the adjoint operator T^*.

Our state space Σ being compact, the Markov operator T is tight. Moreover, T satisfies the Feller property, that is to say that Tf is continuous and bounded whenever f is. We deduce from these two properties the existence of at least one invariant probability measure λ. To reach stronger results, we adopt Canning's key behavioral assumption (Canning, 1992):

> "Agents sometimes make mistakes, choosing an action that is independent of their history. This condition implies that the empirical distribution of outcomes in the model converges to a unique limit distribution."

Allowing agents to make mistakes naturally makes our model more realistic as well. Notice that the word mistake could also encompass innovative strategies or overlapping generation models. At each time step:

- with a probability p, agent n^o i makes a mistake and adopts the first pure strategy with the initial probability $I^{(i)} = p_{0,0}^{(i)}$,

- with a probability $(1-p)$, agent n^o i plays according to the probabilities obtained from the replicator dynamic equations.

At this point, we need to enlarge the space of events to take into consideration possible mistakes. Let T_p denote the Markov operator when agents make mistakes. Using Doeblin's condition, we can show that T_p is quasi-compact.

Moreover, the state $\overline{s} = (A, I^{(1)}, I^{(2)}, I^{(3)}, 0, 0, a^{(3)}, b^{(1)}, b^{(2)}, 0, \frac{b^{(1)}}{a^{(3)}} B_0)$ satisfies the generalized uniqueness criterion: for any integer $k \geq 1$, any state $s \in \Sigma$ and any neighborhood U of $\overline{s}$, there exits an integer n such that $P^{nk}(s, U) > 0$.

By allowing agents to make mistakes, we have given a special role to the state $\overline{s}$ (and to other states of the same form). Indeed, if the composite system is in an arbitrary state s_k at time t_k, then the state $\overline{s}$ will be reached at time t_{k+3} with a strictly positive probability independent of s_k. The proofs of the above two properties rely on this remark.

From the above, we conclude that:

- there exists a unique invariant probability measure λ_p,

- the sequence of probability measures $\{(T_p^*)^n \mu\}$, for any initial probability measure μ, converges to λ_p at a geometric rate in the topology induced by the total variation norm.

In words, the unique invariant probability measure λ_p reflects the long run average (Cesaro sequence: $\frac{1}{n} \sum_{i=0}^{n} (T_p^*)^i \mu$) and the limit behavior for any initial condition. Moreover, when the probability of making a mistake p tends to zero, the equilibrium distribution λ_p converges to an invariant distribution λ for the model without mistakes. Formally, the equilibrium correspondence $E : p \to \lambda_p$ is upper hemi-continuous in the topology of weak convergence at $p = 0$. Mistakes can thus be considered as an equilibrium selection device. However, as Canning (1992) warns us:

> "While mistakes are a refinement, they do not necessarily pick out a unique equilibrium; in some cases the distribution of the mistakes, which actions are chosen if a mistake is made, may affect which equilibrium is selected."

Montin (2004) provides proofs of the properties of this section.

4 Simulations

All simulations have been performed with Matlab.

Simulated stock price fluctuations over one year (250 trading days) are plotted in figure 4. At each time step (for simplicity's sake, $t_k = k$), a random number generator is used to make the composite system collapse to an attainable configuration according to the transition probabilities.
We have used the following set of parameters:

$$\begin{cases} S_0 = B_0 = 1 \\ a_0^{(1)} = a_0^{(2)} = 0 \quad a_0^{(3)} = 6 \quad b_0^{(1)} = 4 \quad b_0^{(2)} = 3 \quad b_0^{(3)} = 0 \\ p = 0.02 \quad \text{(the probability of making a mistake)} \\ r = (1.1)^{\frac{1}{250}} - 1 \quad \text{(the annual riskfree rate of interest is 10\%)} \end{cases}$$

It is quite intuitive to understand that the stock's high volatility comes from its lack of liquidity. A model involving more agents would most probably lead to lower daily average returns and volatilities. The high historic annualized Sharpe ratio (approximately 2.09) gives a second explanation to why agents might like stocks.

Figure 4: stock price fluctuations

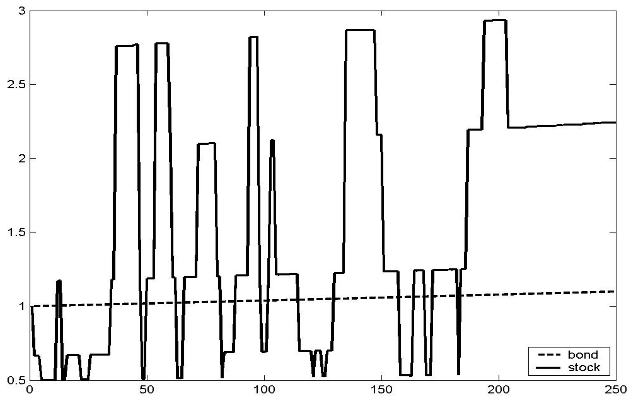

As a simple monetary policy, we now study the impact of a decrease of the riskfree rate of interest r on the limit distribution of the discounted stock price. Recall that the iterates $\{P^n(s,.) = (T_p^*)^n \delta_s\}$ converge to the invariant distribution λ_p at a geometric rate in the topology induced by the total variation norm. When agents make mistakes, there are 32 possible succeeding states s_{k+1} to s_k. Exploring the entire tree to compute the exact distribution after n iterations quickly becomes unrealistic. We thus adopt a frequentist approach. Figure 5 represents the empirical distributions of the discounted stock price after respectively 200 and 250 iterations. Both distributions have been constructed with $15,000$ independent sample paths. To better visualize the impact of r, we have chosen the one-period riskfree rate to be equal to $r = 0.05$. We have also computed the distance induced by the total variation norm between the two distributions ($d \simeq 0.11$) to measure the quality of convergence toward the unique invariant distribution λ_p.

Figure 5: discounted stock price distributions after 200 and 250 iterations

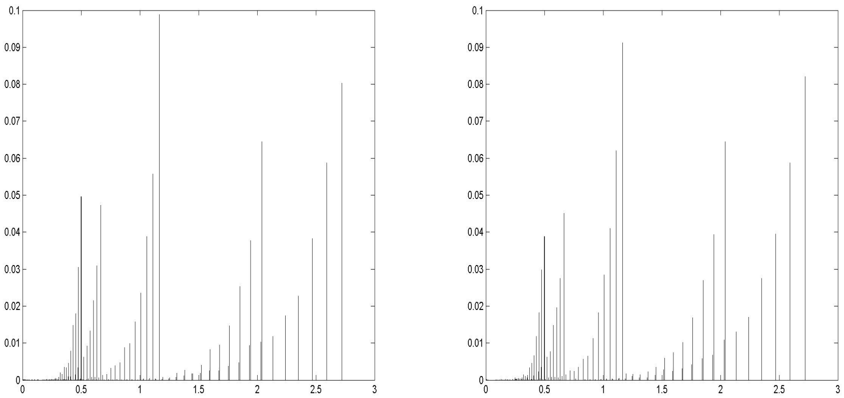

Qualitatively, a decrease of the riskfree rate of interest r has two complementary consequences on the shape of the limit distribution of the discounted stock price.

- As it can be observed in figure 6, positive rates of interest induce the existence of local left-tails shifting the mean of the discounted stock price to lower values. The larger the rate of interest is, the wider the local tails are.

- Intuitively, decreasing the riskfree rate of interest r makes stocks more attractive. From the replicator equations, it is easy to prove that the probability $p_{1,k}^{(i)}$, with which agent n^o i is willing to own stocks at time t_{k+1}, indeed increases when r decreases. It is thus more likely to be in a configuration of the second class (two agents own stocks) for lower values of r. Discounted stock price fluctuations resulting from transitions among configurations of the second class correspond to the right side of the discounted stock price distribution.

Figure 6: discounted stock price distributions after 250 iterations for $r = 0.05$, $r = 0.025$, $r = 0.01$ and $r = 0$

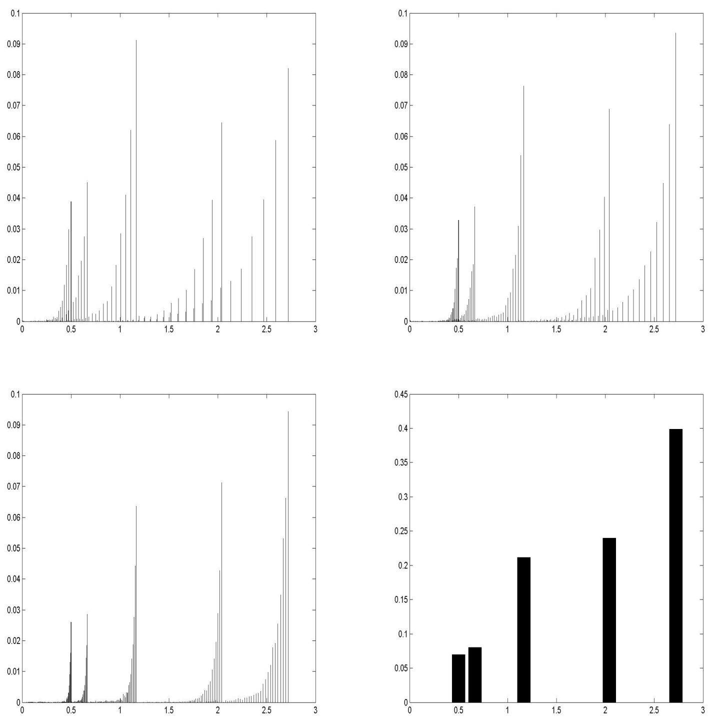

Finally, figure 7 gives the empirical distribution of the one-period stock log returns after 250 iterations: $\ln\left(\frac{S_{250}}{S_{249}}\right)$. The probability that the log return equals zero is approximately 0.6664. To better visualize the tails of the distribution, the $y-$axis stops at the value 0.045. Again, the distribution has been constructed with $15,000$ independent sample paths and the daily risk free rate of interest is 0.05. For comparison, the Gaussian distribution, with the same first two moments, is also represented. It is striking and comforting that the empirical distribution's peak around the mean is higher and narrower and that the tails are fatter (leptokurtic). This is a nice result since more traditional financial models usually fail to consider or explain such real world features (e.g., Pagan, 1996). The skewness and kurtosis are:

$$\text{Skewness} \simeq 1.07 \qquad \text{Kurtosis} \simeq 8.03$$

The above values partially explain why agents might like stocks. The expected net stock return over one period is approximately 0.05, that is to say very similar to the riskfree rate of interest r. But there are states of the world with very high payoffs. Friedman and Savage's (1948) double-inflection utility function could explain the agents' behavior: depending on the level of wealth, an agent is risk-averse (poor) or risk-loving (wealthy).

Figure 7: log returns, Gaussian and hyperbolic distributions

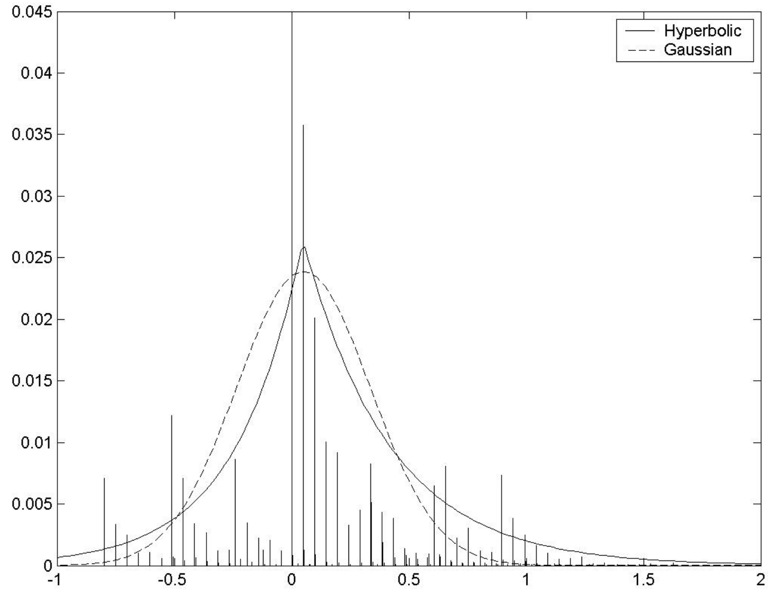

The class of generalized hyperbolic distributions exhibits the observed semi heavy-tails. The subclass of hyperbolic distributions was introduced in finance by Eberlein and Keller (1995). Their Lebesgue density can be expressed as:

$$d_H(x) = \frac{\sqrt{\alpha^2 - \beta^2}}{2\alpha\delta K_1(\delta\sqrt{\alpha^2 - \beta^2})} \exp\left(- \alpha\sqrt{\delta^2 + (x-\mu)^2} + \beta(x-\mu)\right) \quad (34)$$

where K_1 denotes the modified Bessel function of the third kind. The improvement in fitting our empirical log return distribution with an hyperbolic distribution instead of a Gaussian distribution is as well illustrated in figure 7. The parameters have been estimated by maximum likelihood with the free statistical software R (we have also truncated the peak to better fit the tails).

5 Conclusion

Other aspects, such as option pricing, can easily be studied in the frame of our model. If only two agents were present, then our setting would fall under the extensively studied Cox-Ross Rubinstein binomial tree (1979). Indeed, only two events could occur: the agents remain in their respective states or they switch allocations. With three agents, given the stock price S_k at time t_k, there are four possible values for S_{k+1}. With two assets, it is not possible to replicate an arbitrary contingent claim: the market is incomplete. The potential buyers and sellers have different goals and set different prices. We can follow Mel'nikov (1999) for an intuitive definition of the bid and ask prices:

- To hedge the claim's risk, the seller is ready to sell the claim for a price equal to the most inexpensive portfolio that ensures a cashflow at maturity greater than or equal to the claim's cashflow.

- Conversely, the buyer of the claim is ready to short-sell the most expensive portfolio which cashflow at maturity is covered by the claim's cashflow.

The interested reader should refer to Mel'nikov (1999) and Montin (2004) for a more complete analysis with examples.

Most agent-based computational economies heavily rely on simulations. Having adopted a simple representation of financial markets, we have been able to prove theoretical results and gain intuition on complexity economics. Of interest, the limit empirical stock log return distribution presents real world features usually not taken into account by more traditional models. We hope to create more realistic models as by-products of these first steps.

References

[1] Chiu Fan Lee and Neil F. Johnson, 2003. Efficiency and formalism of quantum games. Physical Review A, 67 (2), reference 022311.

[2] David Canning, 1992. Average behavior in learning models. Journal of Economic Theory, 57, 442-472.

[3] Drew Fudenberg and David K. Levine, 1998. The Theory of Learning in Games. The MIT Press.

[4] Larry Samuelson, 1997. Evolutionary Games and Equilibrium Selection. The MIT Press.

[5] Benoît Montin, 2004. A Stock Market Agent-Based Model Using Evolutionary Game Theory and Quantum Mechanical Formalism. PhD thesis, FSU.
www.math.fsu.edu/∼bmontin

[6] John Stachurski. Lagrange stability in economic-systems with positive invariant density. Unpublished.

[7] Carl A. Futia, 1982. Invariant distributions and the limiting behavior of Markovian economic models. Econometrica, 50 (2), 377-408.

[8] Nelson Dunford and Jacob T. Schwartz, 1957. Linear Operators (Part I). Interscience Publishers, New York.

[9] Adrian Pagan, 1996. The econometrics of financial markets. Journal of Empirical Finance, 3 (1), 15-102.

[10] Milton Friedman and Leonard Jimmie Savage, 1948. The utility analysis of choices involving risk. Journal of Political Economy, 56 (4), 279-304.

[11] Ernst Eberlein and Ulrich Keller, 1995. Hyperbolic distributions in finance. Bernoulli, 1, 281-299.

[12] John C. Cox, Stephen A. Ross and Mark Rubinstein, 1979. Option pricing: a simplified approach. Journal of Financial Economics, 7 (3), 229-263.

[13] Alexander V. Mel'nikov, 1999. Financial Markets: Stochastic Analysis and the Pricing of Derivative Securities in: Translations of Mathematical Monographs, Vol. 184, American Mathematical Society.