



BUILDING AN ARTIFICIAL STOCK MARKET POPULATED BY REINFORCEMENT-LEARNING AGENTS

Aleksandras Vytautas Rutkauskas¹, Tomas Ramanauskas²

Vilnius Gediminas Technical University, Saulėtekio al. 11, LT-10223 Vilnius, Lithuania

E-mails: ¹ar@vgtu.lt; ²finansu.katedra@vgtu.lt

Received 10 April 2008; accepted 26 September 2009

Abstract. In this paper we propose an artificial stock market model based on interaction of heterogeneous agents whose forward-looking behaviour is driven by the reinforcement-learning algorithm combined with some evolutionary selection mechanism. We use the model for the analysis of market self-regulation abilities, market efficiency and determinants of emergent properties of the financial market. Distinctive and novel features of the model include strong emphasis on the economic content of individual decision-making, application of the Q-learning algorithm for driving individual behaviour, and rich market setup. Along with that a parallel version of the model is presented, which is mainly based on research of current changes in the market, as well as on search of newly emerged consistent patterns, and which has been repeatedly used for optimal decisions' search experiments in various capital markets.

Keywords: artificial stock market model, market price, agent heterogeneity, stock value.

1. Introduction

In this paper we develop an artificial stock market (ASM) model, which could be used to examine some emergent features of a complex system comprised of a large number of heterogeneous learning agents that interact in a detail-rich and realistically designed environment. This version of the model is not calibrated to empirical data, so at this stage the main aim of this research is to offer, implement and test some new ideas that could lay ground for a robust framework for analysis of financial market processes and their determinants. We believe that the model does offer an interesting framework for the structured analysis of market processes without abstracting from relevant and important features, such as an explicit trading process, regular dividend payouts, trading costs, agent heterogeneity, dissemination of experience, competitive behaviour, agent prevalence and forced exit, etc. Of course, some of these aspects have already been incorporated in existing agent-based financial models. However, the lack of the widely accepted fundament in this area of modelling necessitates the individual

and largely independent approach, which is pursued in this study. Along with that an alternative from this viewpoint decisions management system in capital markets is being analysed, which is based on certain assumptions about continuity of capital market behaviour and on newly formed features, and it allows being efficiently used in various capital markets during global financial crisis conditions.

One of the most interesting features of the artificial stock market modelling is a relatively detailed modelling of decision processes. In our view, advantages of agent-based models for deepening our understanding of real world financial processes can only be fully utilised if a strong emphasis is put on the economic content of the model – individual behaviour and market structure must be based on clear and economically sound principles. Importantly, agents in our model exhibit economically appealing and forward-looking behaviour, which is based on adaptive learning, more specifically, a combination of reinforcement-learning (Watkins' Q-learning) and evolutionary change. To our knowledge, this is one of the first attempts to incorporate Q-learning algorithm into the ASM models.

By conducting simulation experiments in this model, we aim to address some specific questions, such as market self-regulation abilities, the congruence between the market price of the stock and its fundamentals (the market efficiency issue), importance of intelligent individual behaviour and interaction at the population level for market efficiency and functioning, and relationship between stock prices and market liquidity.

2. Description of the ASM model

The ASM research area is relatively new but there is a growing body of literature on the subject. There is a clear lack of the comprehensive literature review and classification of existing models. Some popular models and ASM modelling principles presented in LeBaron (2006), Samanidou *et al.* (2007) have a review of some agent-based financial models, with the emphasis on econophysics. Interaction of heterogeneous agents lies at the heart of ASM models, which leads to complex systemic behaviour and emergent systemic properties. There are two broad classes of ASM models, namely, models based on agents' hard-wired behavioural rules (see, e.g. Kim and Markowitz (1989), Sethi and Franke (1995), Lux (1995)) and models supporting systemic adaptation. The most prominent example of the latter category is the Santa Fe ASM model developed by Arthur *et al.* (1997); also see, e.g. Beltratti and Margarita (1992), Lettau (1997), LeBaron (2000), Tay and Linn (2001). See Ramanauskas (2009) for a general discussion about agent-based financial modelling and the above-mentioned models. In many models systemic adaptation is usually warranted by evolutionary algorithms, whereas individual agents' behaviour is very stylised and based on economic consideration directly. In contrast, in modelling financial market processes we put a strong emphasis on individual behaviour and economic reasoning.

The present ASM model does not fully abstract from many important features of real financial markets that are usually excluded both from standard financial models and other ASMs. For example, just like in the real world financial markets, agents in this ASM do not know the "true model" but try instead to adapt in the highly uncertain environment, they exhibit bounded rationality, non-myopic forward-looking behaviour, as well as diversity in experience and skill levels; the trading process is quite realistic and detailed; dividends are paid out in discrete time intervals and the importance of dividends as a fundamental force driving stock prices is explicitly recognised. The proposed

ASM model embodies some new ideas about financial market modelling and provides interesting generative explanation of prolonged periods of over- and under-valuation. In the following section we present the architecture of the artificial stock market in detail.

2.1. General market setting and model's main building blocks

The artificial stock market is populated by a large number of heterogeneous reinforcement-learning investors. Investors differ in their financial holdings, expectations regarding dividend prospects or fundamental stock value. This ensures diverse investor behaviour even though the basic principles governing experience accumulation are the same across population. The very basic description of agents' behavioural principles can be described as follows. All agents forecast an exogenously given, unknown dividend process and base their estimates of the fundamental stock value on dividend prospects. These estimates are intelligently adjusted to attain immediate reservation prices. Agents explore the environment and accumulate the experience with the aim of maximising long-term returns on their investment portfolios but there are no optimality guaranties in the context of high uncertainty and complex interaction of agents.

As usual in financial market modelling, the modelled financial market is very simple. Only one, dividend-paying stock (stock index) is traded on the market. Dividends are generated by an exogenous stochastic process unknown to the agents, and they are paid out at regular intervals. The number of trading rounds between dividend payouts can be set arbitrarily, which enables interpretation of a trading round as a day, a week, a month, etc. Paid out dividends and funds needed for liquidity purposes are held in private bank accounts and earn constant interest rates, whereas liquidity exceeding some arbitrary threshold is simply removed from the system (e.g., consumed). Borrowing is not allowed. Initially agents are endowed with arbitrary stock and cash holdings, and subsequently in every trading round each of them may submit a limit order to buy or sell *one unit* of stock, provided, of course, that financial constraints are non-binding. Trading takes place via the centralised exchange.

For the ease of detailed model exposition, it is useful to break the model into a set of economically meaningful processes, though some of them are interrelated in complex ways. The general structure of the model is laid out in Fig. 1. We will discuss these logical building blocks in the following subsections.

Forming private forecasts of exogenously generated dividends Based on: <ul style="list-style-type: none"> • Exponential moving average • Adjustment as a result of reinforcement-learning (agents seek to minimise forecast errors)
Making individual estimates of fundamental stock value and its reservation price Based on: <ul style="list-style-type: none"> • Discounted expected dividend flows • Adjustment as a result of reinforcement-learning (agents seek to maximise portfolio returns)
Making individual trading decisions Based on: <ul style="list-style-type: none"> • Private estimates of fundamentals, • Maximisation of expected individual wealth at the end of a trading period • Publicly announced estimated probabilities of successful trades for given prices
Carrying out trades via the centralised exchange and collecting trading statistics Based on: <ul style="list-style-type: none"> • Double auction system • Simultaneous submission of trade orders and random queuing of individual orders
Learning to forecast dividends and learning about fundamental stock value Based on: <ul style="list-style-type: none"> • Standard Q-learning with linear gradient-descent approximation
Augmenting learning processes by specific interaction among agents (optional) Based on: <ul style="list-style-type: none"> • Successful strategy imitation • Evolutionary selection and resultant prevalence of successful investment strategies • Noise trading behaviour

Fig. 1. Main building blocks of the ASM model

2.2. Forecasting dividends

Expected company earnings and dividend payouts are the main fundamental determinants of the intrinsic stock value. We assume that all agents make their private forecasts of dividend dynamics. Dividend flows generated by an unknown, potentially non-stationary data generating process are specified by a modeller. The only information, upon which agents can base their forecasts, is past realisation of dividends, and agents know nothing about stationarity of the data generating process. Hence, they are assumed to form adaptive expectations, augmented with the reinforcement-learning calibration. We also allow for possibility to improve a given agent’s forecasting ability by probabilistic imitation of more successful individuals’ behaviour (see Section 2.5 for more on this).

Agents start with finding some basic reference points for their dividend forecasts. The exponentially weighted moving average (EWMA) of realised dividend payouts can be calculated as follows:

$$d_{i,y}^{EWMA} = \lambda_1 \cdot d_y + (1 - \lambda_1) d_{i,y-1}^{EWMA}. \quad (1)$$

Here d_y denotes dividends paid out in period y (year) and λ_1 is the arbitrary smoothing factor (the same for all agents), which is a real number between 0 and 1. The subscript i on the averaged dividends in equa-

tion (1) is to indicate that they vary across individual agents. The differences arise due to different arbitrarily chosen initial values but over time, however, these exponential averages converge to each other. Also note that dividend payouts can be arbitrarily less frequent than stock trading rounds, e.g. if one trading period equals one month, dividends may be scheduled to be paid out every twelve periods and in equation (1) one time unit would be one year.

Exponential moving averages would clearly be unacceptable estimates of future dividends in a general case. Hence, their function in this model is twofold. First, they provide a basis for further “intelligent” refinement of dividend forecasts, i.e. these moving averages are multiplied by some adjustment factors calibrated in the process of the reinforcement learning. And second, forecasting dividends relative to their moving averages, as opposed to forecasting dividend levels directly, makes forecasting environment more stationary, which facilitates the reinforcement-learning task.

The n -period dividend forecast is given by the following equation:

$$E(d_{i,y+n}) = d_{i,y}^{EWMA} \cdot a_{i,y}^{div}, \quad (2)$$

where $a_{i,y}$ is agent i ’s dividend adjustment factor. These adjustment factors are gradually changed as agents explore and exploit their accumulated experience, with the long-term aim to minimize squared forecast errors. The detailed description of the reinforcement-learning procedure is provided in Section 2.5 and Appendix 1. Individual forecasts for periods $y + 1, \dots, y + n$ formed in periods $y - n + 1, \dots, y$, respectively, are stored in the program and used for determining individual estimates of the fundamental stock value.

2.3. Estimating fundamental stock value and reservation prices

Quite similarly to the dividend forecasting procedure, agents’ estimation of the intrinsic stock value is a two-stage process. It embraces formation of initial estimates of the fundamental value, based on discounted dividend flows, and ensuing intelligent adjustment grounded on agents’ interaction with environment. We refer to this refined fundamental value as the reservation price.

The initial evaluation of the future dividend flows is a simple discounting exercise. To calculate the present value of expected dividend stream, the constant interest rate is used as the discount factor. Moreover, beyond the forecast horizon dividends are assumed to remain constant. Under these assumptions, individual

estimates of the present value of expected dividend flows are:

$$v_{i,y}^{fund} = d_y + E \left(\frac{d_{i,y+1}}{1+\bar{r}} + \dots + \frac{d_{i,y+n}}{(1+\bar{r})^n} + \frac{d_{i,y+n}/\bar{r}}{(1+\bar{r})^{n+1}} \right), \quad (3)$$

where \bar{r} is the constant interest rate. The last term in this equation is simply the discounted value of the infinite sum of steady financial inflows. These present value estimates are subject to further refinement.

To avoid excessive volatility of the estimates of the discounted value of dividend stream, they are again smoothed by calculating the exponentially weighted moving averages:

$$v_{i,y}^{EWMA} = \lambda_2 \cdot v_{i,y}^{fund} + (1 - \lambda_2)v_{i,y-1}^{EWMA}. \quad (4)$$

The role of these averages is very similar to that of the averaged dividends in the dividend forecasting process, namely, to provide some background for the reinforcement-learning procedure and (partially) stationarise the environment in which agents try to adapt.

The second stage in the estimation of the individual reservation prices of the stock is calibration based on the reinforcement-learning procedure. For this we have to switch to the different time frame (in the base version of the model it is assumed that dividends are paid out annually, whereas agents can trade once per month). In a given trading round t , individual reservation prices $v_{i,t}^{reserve}$ are obtained from equation (4) by multiplying exponentially smoothed estimates of fundamental value by individual price adjustment factors, $a_{i,t}^p$:

$$v_{i,t}^{reserve} = v_{i,t}^{EWMA} \cdot a_{i,t}^p. \quad (5)$$

In this context the individual reservation price is understood as an agent's subjective assessment of the stock's intrinsic value that prompts immediate agent's response (to buy or sell the security).

2.4. Making individual trading decisions

Having formed their individual beliefs about the fundamental value of the stock price, agents have to make specific portfolio rebalancing decisions. In principle, they weigh their own assessment of the stock against market perceptions and make orders to buy (sell) one unit of the underpriced (overpriced) stock at the price that is expected to maximise their wealth at the end of the trading period. We give a more detailed description of these processes below.

The individual reservation price reflects what investors think the stock price should be worth. If the last period's average market price p_{t-1} is less than agent i 's reservation price today, it is willing to buy stock

and pay at most $v_{i,t}^{reserve}$. Conversely, if the prevailing market price is higher than the agent's perceived fundamental, it is willing to sell it at $v_{i,t}^{reserve}$ or higher price. So its decision rule is like this:

If $v_{i,t}^{reserve} > p_{t-1}$ and $m_{i,t}^0$ is sufficient \rightarrow submit limit order to buy 1 share at price $p_{i,t}^q$

if $v_{i,t}^{reserve} < p_{t-1}$ and $h_{i,t}^0 > 0 \rightarrow$ submit limit order to sell 1 share at price $p_{i,t}^q$

otherwise, make no order.

Here $h_{i,t}^0$ and $m_{i,t}^0$ denote, respectively, agent i 's stock holdings (i.e. number of owned shares) and cash balance at the beginning of a trading round, $p_{i,t}^q$ is the quoted price to be determined below.

We would not expect real world investors to make orders to buy or sell the stock precisely at reservation prices because in that case they would miss potentially profitable asset allocation opportunities. The real world investor whose perception of the stock value considerably differs from the average market opinion is likely to take advantage of market liquidity and make an order to trade at a price close to the prevailing market price rather than to his own reservation price. But what price would it be? There is no answer in the theory. The first obvious step, implemented in the model, is to allow limit orders, i.e. orders to trade the security at a specified *or* better price. Given the complexity of the agents' interaction, the optimal pricing solution generally cannot be found. Thus, we proceed in the following, intuitively appealing way: (i) we determine the possible price quote grid around the prevailing market price (i.e. determine tick sizes and possible price fluctuation bands), (ii) estimate aggregate supply and demand schedules, (iii) compute each individual's expected end-of-period wealth for every possible trading price and (iv) allow agents to make trading decisions that maximise their expected end-of-period wealth.

Agents, of course, aim at getting most favourable prices for their trades but they must take into account the fact that better bid or ask prices are generally associated with smaller probabilities of successful trades. The assumption that each agent is allowed to trade only one unit of stock in a given trading round has a very useful implication in this context – the probabilities of successful trades at all possible prices faced by a buyer and a seller can be loosely interpreted as the supply and demand schedules, respectively. So we further assume that these supply and demand schedules are estimated by the exchange institution from past trading data and constitute public knowledge.

Estimated probabilities of successful trades at given (relative) price quotes are computed as follows. Sim-

ply put, these estimated probabilities should indicate chances of successful trading at prices that are “high” or “low” relative to the prevailing market price (i.e. last period’s average price). So the probability of the successful trade for a given price quote (relative to the benchmark price) is calculated from the past trading rounds as a fraction of successfully filled buy (sell) orders out of all submitted orders to buy (sell) at that price. Unfortunately, due to computational constraints the number of agents and successful trades is not sufficiently high to obtain reliable estimated probabilities in this straightforward way. For this reason we employ the following three-step procedure:

- i) estimates of probabilities of successful buy and sell orders for every price quote are smoothed over time by computing exponential moving averages;
- ii) if there are no orders to buy or sell at a given price at time t , the exponential moving average estimates of successful trade probabilities are left unchanged from the $t-1$ period;
- iii) the scattered estimates are fitted to a simple cross-sectional regression line (with its values restricted to lie in the interval between 0 and 1) to ensure that the sets of successful trade probabilities retain meaningful economic properties.

As a result, we get a nice upward-sloping line, which represents probabilities of successful buy orders for each possible price quote, and a downward-sloping line for the sell orders case. Fig. 2 shows a typical example of estimated probabilities of successful buying and selling one unit of stock at all possible prices (last period’s average price set equal to 25 in this relative pricing grid). This particular example reflects an upward-trending market, in which agents reckon they have higher chances (estimated at around 60%) of sell-

ing the stock than buying it (estimated at around 40%) at the last period’s average price.

At this stage agents have all the components needed to choose prices that give them highest expected wealth at the end of the trading round. First, agent i estimates its expected end-of-period stock holdings (i.e. the number of shares) for each possible price quote j :

$$E(h_{i,j,t}^1) = h_{i,t}^0 + E(q_{i,j,t}) \cdot b_i \text{ for all } j. \quad (6)$$

Here $E(q_{i,j,t})$ denotes expected number of shares to be bought or sold by agent i at any quotable price j (as was explained above, these numbers lie in the closed interval between 0 and 1). The indicator variable b_i takes value of 1 if the agent is willing to buy the stock or -1 if it is willing to sell the stock.

Similarly, agent i ’s expected end-of-period cash holdings for each possible price quote j are:

$$E(m_{i,j,t}^1) = m_{i,t}^0 + E(q_{i,j,t}) \cdot x_{j,t} \cdot (-b_i - c) + E(h_{i,j,t}^1) \cdot E(d_{i,t}) \text{ for all } j. \quad (7)$$

Here $x_{j,t}$ denotes possible price quote j , c is the fractional trading cost and $E(d_{i,t})$ denotes the expected dividends, which are to be paid out following the trading round (this term equals zero in between the dividend payout periods). It is important to note here that the interest on spare cash funds is paid, as well as excess liquidity (cash holdings above some prespecified amount needed for trading) is taken away, at the beginning of the trading period. All of this is reflected in $m_{i,t}^0$. Dividends are paid out for those agents that hold stocks after the trading round, as can be seen from equation (7).

Finally, agent i ’s expected end-of-period stock holdings are valued at individual reservation price and each agent calculates its expected end-of-period wealth for every possible price quote:

$$E(w_{i,j,t}^1) = E(h_{i,j,t}^1) \cdot v_{i,t}^{reserve} + E(m_{i,j,t}^1). \quad (8)$$

Hence, agent i ’s quoted price, p_i^q , is the price that is associated with the highest expected wealth at the end of the trading round:

$$p_i^q = \arg \max_{x_i} E(\bar{w}_{i,t}^1). \quad (9)$$

If several price quotes result in the same expected wealth, the agent chooses randomly among them. It is also important to note that in the process of the reinforcement learning, agents are occasionally forced to take exploratory actions. In those cases exploring agents choose prices from the quote grid in a random manner.

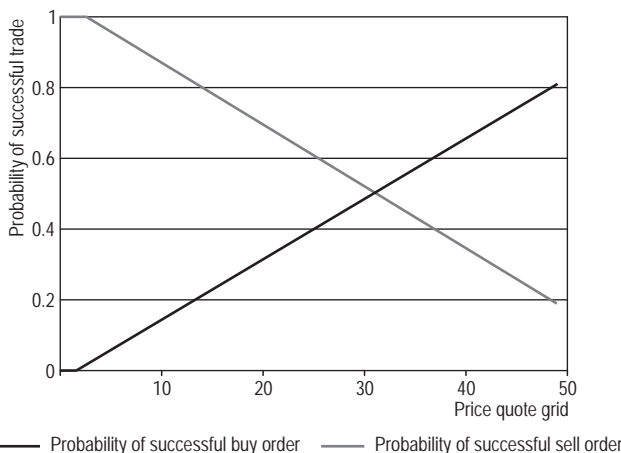


Fig. 2. Typical estimated demand and supply schedules in an upward-moving market

Market price determination and actual trading take place on the centralised stock exchange. The trading mechanism basically is the double auction system, in which both buyers and sellers contemporaneously submit their competitive orders to implement their trades. Agents are assumed to have no knowledge of individual market participants' submitted orders.

In this model the order book mechanism works as follows. Prior to a trading round, all agents' trade orders are queued randomly and then each of them undergoes the processing procedure. During this procedure, for an order that is being processed all earlier-queued orders are scanned in search for the most favourable matching (opposite) order. If such an order is found (a tie among several equally good orders is broken arbitrarily), the trade is filled at the average of the bid and ask price. Otherwise, the order remains open until it makes a match for other subsequently processed orders or until the end of the trading period, when it is closed as an unexecuted order. Following the trading round, all agents' cash and securities accounts are updated accordingly.

The centralised stock exchange also produces a number of trading statistics, both for analytical and computational purposes. These statistics include the market price, trading volumes and volatility measures. The market price in a given trading period is calculated as the average traded price. As was mentioned before, it is crucially important for making further trading decisions and it serves as the reference value in the subsequent trading round.

2.5. Learning and systemic adaptation in the model

We assume that agents' behaviour is driven by reinforcement-learning since these learning algorithms borrowed from the machine learning literature seem to be conceptually suitable for modelling investor's behaviour. Agents take actions in the uncertain environment and obtain immediate rewards associated with these (and possibly previous) actions. A specific learning algorithm allows agents to adjust their action policies in pursuit of highest long-term rewards. It is a very desirable feature of any financial model that agents strive for strategic, as opposed to myopic, behaviour. The reinforcement-learning agents do just that. On the other hand, it is the immense complexity of investors' interaction, both in real world financial markets and in the model, that dramatically limits agents' abilities to actually achieve optimal investment policies if not makes the optimal investment behaviour outright impossible. In our model we use a popular reinforcement-learning

algorithm, also known as the Q-learning, which was initially proposed by Watkins (1989). It is the temporal difference learning based on the step-wise update (or back-up) of the action-value function and associated adjustment of behavioural policies (a more detailed exposition of basic Q-learning principles is given in Appendix 1). The principal back-up rule is closely related to Bellman optimality property and takes the following form:

$$Q(s_t, a_t) \leftarrow (1 - \alpha) \cdot \underbrace{Q(s_t, a_t)}_{\text{Old estimate of } Q(s_t, a_t)} + \underbrace{\alpha(r_{t+1} + \gamma \max_a Q(s_{t+1}, a))}_{\text{New estimate of } Q(s_t, a_t)}. \quad (10)$$

Here s_t denotes the state of environment, a_t is the action taken in period t and r_{t+1} is the immediate reward associated with action a_t (and possibly earlier actions). Parameter α is known as the learning rate and γ is the discount rate of future rewards. Function $Q(s_t, a_t)$ is usually referred to as the action-value function (or Q-function) and it basically shows the value of taking action a_t in state s_t under behavioural policy π . More specifically, the action-value function is the expected cumulative reward conditional on the current state, action and pursued behavioural policy.

However, the so-called "curse of dimensionality" implies that the straightforward implementation of the basic version of this algorithm is rarely possible in complicated environments. Following the standard practice, we apply the Q-learning algorithm with gradient-descent approximation. Here we only describe specific variables that are used in the Q-learning algorithm.

As was mentioned before, there are two instances of individual agent learning in the model: learning to forecast dividends and learning to adjust perceived fundamentals. In the dividend forecasting case agent i learns to adjust the dividend adjustment factor, $a_{i,t}^{div}$ (see equation (2)). In each state there are three possible actions – the agent can increase the dividend adjustment factor by a small proportion specified by the modeller, decrease it by the same amount or leave it unchanged.

Due to the complex nature of environment, the state of the world – as perceived by investor i – must be approximated, and it is described by a vector of so-called state features, $\vec{\phi}_s$. We choose four state features that are indicative of the reinforcement learner's "location" in the environment and summarize some properties of the dividend-generating process, which can provide basis for successful forecasting. These features include

the size of the dividend adjustment factor, relative deviation of current dividend from its EWMA (compared to the standard deviation), the square of this deviation (to allow for nonlinear relation with forecasts) and the size of the current dividend relative to the EWMA.

The forecast decision is taken at time y and the actual dividend realisation is known at forecast horizon $y + n$. Then agent i gets the reward, which is the negative of the squared forecast error:

$$r_{i,y+n}^d = -\left(d_{y+n} - E_y(d_{i,y+n})\right)^2. \quad (11)$$

Hence, the agent is punished for the forecasting errors. The learning process is augmented by modeller-imposed constraints on dividend forecasts. The forecast is not allowed to deviate by more than a prespecified threshold (e.g. 30%) from the current level of dividends. In that case, the agent gets extra-punishment and the dividend forecast is forced to be marginally closer to the current dividend level. Once the agent observes the resultant state, i.e. the actual dividend realisation, it updates its behavioural policy according to the Q-learning procedure.

In the case of the individual stock value estimation, agent i can also take one of three actions: fractionally increase or decrease the price adjustment factor, $a_{i,t}^p$ (see equation (5)), or leave it unchanged. Analogously to the dividend forecasting case, the four state features are the price adjustment factor, the stock price deviation from its exponential time-average (this difference is divided by the standard deviation), the square of this deviation and the current stock price divided by the weighted time-average.

The agent observes the state of the world and acts according to the pursued policy. After the trading round, the agent observes trading results and the resultant state of the world, which enables the agent to update its policies according to the usual Q-learning procedure. In this model, the basic immediate reward, $r_{i,t+1}^p$, is simply the log-return on the agent's portfolio:

$$r_{i,t+1}^p = \ln\left(h_{i,t}^1 p_t + m_{i,t}^1 (1 + \bar{r}^{monthly})\right) - \ln\left(h_{i,t}^0 p_{t-1} + m_{i,t}^0\right). \quad (12)$$

Recall that p_t denotes the market price following a trading round in time t and $\bar{r}^{monthly}$ is a one-period return on bank account. In order to ensure more efficient learning – just like in the case of dividend learning – constraints are imposed on the magnitude of price adjustment factors, and additional penalties are invoked if these constraints become binding.

The chosen specification of the reward function implies that the reinforcement-learning agents try to learn

to organise their behaviour so that they maximise long-term returns on their investments. We could interpret agents in this model as professional fund managers that care about maximising clients' wealth, seek best long-term performance among peers and shun under-performance. They need not be risk-averse, as is conventionally assumed about individual consumption-smoothing investors. Indeed, recent evidence from extremely turbulent financial markets shows that it might well be quite the opposite – in some cases excessive risk-taking might generate superior performance for a prolonged period of time, which in turn generates solid growth in fee income during that time. In addition, it should be noted that in the model an agent's attitude toward risk is determined not only by its reward function but also by evolutionary selection and other systemic adaptation.

The model allows for optional alteration of agent behaviour via sharing private trading experience, competitive evolutionary selection and noise trading behaviour. These options help enhance realism of the artificial stock market and arguably augment the reinforcement-learning procedure by removing clearly dominated trading policies implemented by individual agents and by strengthening competition among them.

In our model, dissemination of agents' experience is very stylised. At the end of each period agents are randomly matched in pairs. In every pair, agents' long-term performance measures, which are cumulative past rewards, are compared to each other. If the difference between matched agents' performance measures is sufficiently large (the threshold level is allowed to fluctuate randomly to reflect the random nature of knowledge dissemination), the worse-performing agent simply replicates the more successful agent's experience.

Evolutionary selection is another available option in the present ASM. It assumes bankruptcy of worst-performing agents and their replacement with best-performers. So agents, whose performance relative to the benchmark (which is the average agents' performance) falls below a modeller-specified threshold, go bankrupt. Their place is taken over by best-performers, which then are forced to split so that the number of agents remains constant. This has a natural interpretation: inferior fund managers are forced out of the market as unsatisfied clients bring their wealth over to best-performing funds and the latter then have to split for regulatory or any other reasons. Successful agents are given substantial extra rewards in the event of the split, to encourage their performance.

Finally, the model allows for noise trading behaviour. Unlike in the evolutionary selection, the worst-performers are not replaced by most successful agents. Rather, they scrap their prior learning experience and, as a result, start learning from scratch.

3. Simulation results

Like the vast majority of other ASM models, the current model is based on a large number of parameters, and it is very difficult to calibrate the model to match empirical data. At this stage of the model development we do not attempt to do that. Instead, we assign reasonable and, where possible, conventional values to the parameters and assume very simple forms of dividend-generating processes. This enables us to determine the approximate fundamental stock value dynamics and study how the market stock price, determined by the complex system of interacting heterogeneous agents, fares in relation to stock price fundamentals. Even though the model is not calibrated to the market data, model results can offer qualitative insights about market self-regulation, efficiency and other aspects of market functioning. In this section we examine these issues in more detail and report some of the more interesting simulation results.

The simulation procedure is implemented by performing batches of model runs. Each run consists of 20,000 trading rounds (about 1667 years). Batches of ten runs repeated under identical parameter settings are used to generate essential data and statistics that are in turn used for analysis and generalisation. In every run, the first 5,000 trading rounds – as the learning initiation phase – are excluded from the calculation of the descriptive statistics. The simulation concentrates on altering features of the reinforcement-learning, interaction among agents and dividend-generating processes in an attempt to understand relative importance of intelligent individual behaviour, market setting and population-level changes for the aggregate market behaviour. Other model parameters are kept unchanged.

Dividends are assumed to fluctuate around an exponential trend and their volatility is proportional to the dividend level. The role of the trend is to necessitate the intelligent adjustment of dividend estimates, as forecasts based on exponentially-weighted moving averages would be clearly biased. Large dividend growth rates can only be sustained over relatively short time horizons, and hence in our very long-term model we have to choose very low dividend growth rates (e.g. 0.15% per year). We also examine deterministic constant dividends, as a special case.

The primary question addressed in most ASM models is the market efficiency issue. Here efficiency is loosely interpreted as the congruence between the stock market price and its fundamentals. In the current setting it is not possible to know the right theoretical stock price, so we basically want to compare the market stock price with risk-neutral estimates of fundamentals.

Let us start with the examination of agents' ability to forecast dividends. Since dividends are driven by very simple data generating processes, it is not surprising that in the model version with enabled both reinforcement-learning and evolutionary selection, agents are able to form very precise forecasts. The average dividend forecast error for this model specification is -0.1% , while the average absolute forecast error again amounts to 0.4% . To assess the actual importance of the reinforcement-learning behaviour for dividend forecasting, simulation batches with disabled reinforcement-learning are run (Experiment 3). In these runs agents neither learn to forecast dividends, nor try to optimise their portfolios, as their commensurate reinforcement rewards $r_{i,t+n}^d$ and $r_{i,t+1}^p$ are set to zero. In this case, the average forecast bias considerably increases to -0.8% and the average absolute error stands at 1.4% . In this no-learning case the average percentage of agents hitting the modeller-imposed dividend forecast bounds increases significantly, as compared to the enabled learning case. In other words, learning agents are able to effectively form "reasonable" forecasts, while non-learning agents are simply forced to remain within prespecified boundaries but perform much worse, taken on individual basis. This leads us to a very natural conclusion that in the dividend forecasting process intelligent adaptation matters.

As the next step of our analysis we examine dynamics of the market price in relation to the fundamentals. In Experiment 1 fundamentals anchor the stock price dynamics to some extent, and the market price fluctuates in the vicinity of the perceived fundamental value. The average percentage bias of market price from the fundamentals is low and stands at -1.6% . Nevertheless, the valuation errors are clearly autocorrelated – due to the market inertia and prevailing expectations, the stock price may be above or below risk-neutral fundamentals for extensive periods of time. For instance, runs of uninterrupted overvaluation stretch on average for 44 trading periods and an average length of undervaluation runs is 60 periods. By the same token, average market price deviations from the fundamental valuation are large relative to the price volatility. The enabled evolutionary selection option in the model ensures relatively even wealth distribution among agents

and each trading period active agents (i.e. agents that have sufficient funds and/or stock holdings to trade constitute on average 89.7% of total population). Finally, the average fraction of agents whose adjusted fundamental valuations (reservation prices) fall out of modeller-imposed “reasonable” bounds is very low and stands on average at 0.1% of total population in a trading round.

It turns out that the above results strongly depend on the evolutionary competition assumption. It suffices to disable the evolutionary selection (Experiment 2), and the average percentage stock price bias from the fundamentals boosts to 5.9% along with a dramatic increase in average overvaluation runs to 406. By the end of a simulation run the number of inactive agents per trading round increases to 70–80%, and wealth naturally concentrates in the hands of remaining 20–30% of agents. There are some possible explanations to this overvaluation and wealth concentration. Such overvaluation can be to some extent associated with the model’s feature that excess liquidity is simply taken away from the market, which means that the agents that tend to sell their stock holdings are more likely to consume their money and become inactive. In other words, those agents that highly value the stock tend to dominate in the market. Another interpretation is that worse performing agents are simply driven out of the market. Moreover, a diminishing number of active participants and a smaller degree of competition allows agents to concert their portfolio rebalancing actions in such a way that the market price is driven up, which leads to larger unrealised returns and thereby stronger reinforcement for the remaining active players. These results make sense from the real world perspective. The largest mass of investors want stock prices to be as high as possible (though possibly still compatible with fundamentals), and it is not in their direct interest to have prices that match fundamentals precisely.

We also perform simulations to examine market’s self-regulation ability. In particular, we want to know whether economic forces are strong enough to bring the market to the true fundamentals if they systematically differ from average perceived fundamentals. For this purpose, we introduce an arbitrary upward bias to the estimates of the fundamental value by adding an arbitrary term in equation (3). Then simulation runs are implemented for different model settings, with or without reinforcement-learning. It turns out that the market is not able to find the true risk-neutral fundamentals. In the case of no-learning, stock prices tend to slowly grow larger than the perceived fundamentals. In

the case of enabled reinforcement-learning, agents tend to stick to the perceived fundamentals, and the market price fluctuates around them as a result.

The above results confirm that the market self-regulation mechanism in this model is weak. We do not find evidence of agents adjusting their perceived fundamentals so that the market price gets in line with modeller-imposed fundamentals or, say, the usually assumed risk-averse behaviour. On the other hand, it is not surprising. Well-known puzzles of empirical finance and recent mega-bubbles suggest that, after all, markets may not be tracking fundamentals so closely. It can be the case that markets exhibit so strong inertia that even fundamentally correct investment strategies pay out only in too distant future and may not be applied successfully or act as the market’s self-regulating force. The obtained results suggest that (not necessarily objectively founded) market beliefs of what an asset is worth are a very important constituency of its market price. Last but not least, we want to examine the relationship between the market price fluctuations and the financial market liquidity. This experiment also helps to shed light on the reasons for a relatively loose connection between the market price and fundamentals. In this simulation run, the standard model version with reinforcement-learning and evolutionary selection is used, while dividends are assumed to be deterministic and constant. It is notable that even in this environment market price fluctuations remain significant and trading does not stop. The clue to understanding this excess volatility may be the positive relationship between market liquidity and the stock price. Since unnecessary liquidity at an individual level is removed from the system, overall liquidity fluctuates in a haphazard way. Increases in market liquidity bolster solvent demand for the stock and lift its price. As can be seen from Fig. 3 liquidity growth spikes are associated with strong price increases. The linear correlation between growth of money balances and stock price growth is found to be 0.32.

It should be noted that the latter experiment is devised so as to ensure that positive relationship between stock returns and investors’ cash holdings is not linked to fluctuations in dividend payouts. This allows us to conclude that liquidity fluctuations affect the asset price in this case, and not vice versa. The evidence that market liquidity changes can move markets is very important for understanding the way liquidity crises, credit booms and busts (deleveraging), portfolio reallocations between asset classes and other exogenous factors may affect stock markets.

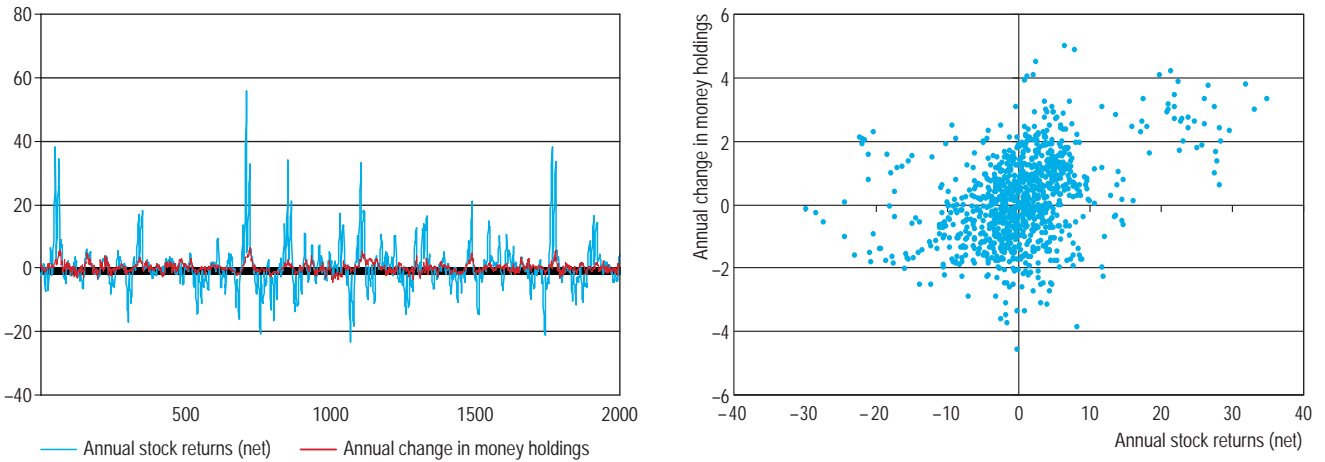


Fig. 3. Typical relationship between stock returns and liquidity in a constant dividend case

4. Short presentation of parallel decisions management system in capital markets

In this paper using the definition of parallel decisions management system in capital and exchange markets we mean the so-called double trump model (Rutkauskas 2008; Rutkauskas *et al.* 2008), which at first was designed for decisions management in exchange markets, and later it was repeatedly used in various capital markets. The description of double trump model, its development and possibilities of application for decisions management in exchange markets can be found in (Rutkauskas 2005; Rutkauskas 2006; Rutkauskas, Stasytė 2006).

The link of parallel system with general title of this paper “Building an artificial stock market populated by reinforcement-learning agents” and ability of market participants and the market itself to match with consistent patterns of market behaviour and decisions management, disclosed by the science, could be described by the following circumstances:

- in financial instability circumstances capital markets sophisticatedly, but cognizably change the supply of possibilities for investor, which is fully enough described by possibilities’ efficiency, riskiness and reliability;
- investor – individual, institutional or criterial has a possibility to perceive the adaptation principles and means of his utility function in changing behaviour of the capital market;
- perception of decisions management strategies and criteria interdependencies becomes a presumption and guarantee of successful investing;

- decisions are made in almost fully artificial market, which is commensurated with real market data only by the core parameters. The space of stochastic processes is an adequate enough real market model.

The possibilities of parallel decisions management system will be illustrated by its application for achievement of as high as possible growth of invested capital during the analysed period: 2007.01.02–2009.04.09, which also includes the most severe periods of global financial crisis. In general, using this system a broad monitoring is organized, which includes about 30 various global markets, for the search of favourable decisions for capital growth. The search of favourable decisions was performed selecting the so-called pseudo-scenario, when a part of historical data – in this case 40 days from the beginning of 2007.01.01 – is accepted as “real” historical data, and next data is treated as “forthcoming”, and with regard to the latter forecasting is being performed and portfolio rebalancing decisions are being made. Next, the results of the experiment in five countries’ markets will be presented: UK, France, Germany, Sweden and USA. In *a* sections of 4–8 Figures we see how a unit of invested capital was changing during the analysed period, and in *b* sections we see the change of real prices of the 6 stocks in portfolio.

The growth of initial invested capital is obtained with the help of optimization of (rebalancing) portfolio structure. While rebalancing the portfolio, a fee of three basic percentile points was imposed in case of buying as well as selling a stock. This amounts to nearly 30% of the whole invested capital and more than 50% of gross capital increase.





Fig. 4. The results of the experiment in UK





Fig. 5. The results of the experiment in Germany





Fig. 6. The results of the experiment in Sweden

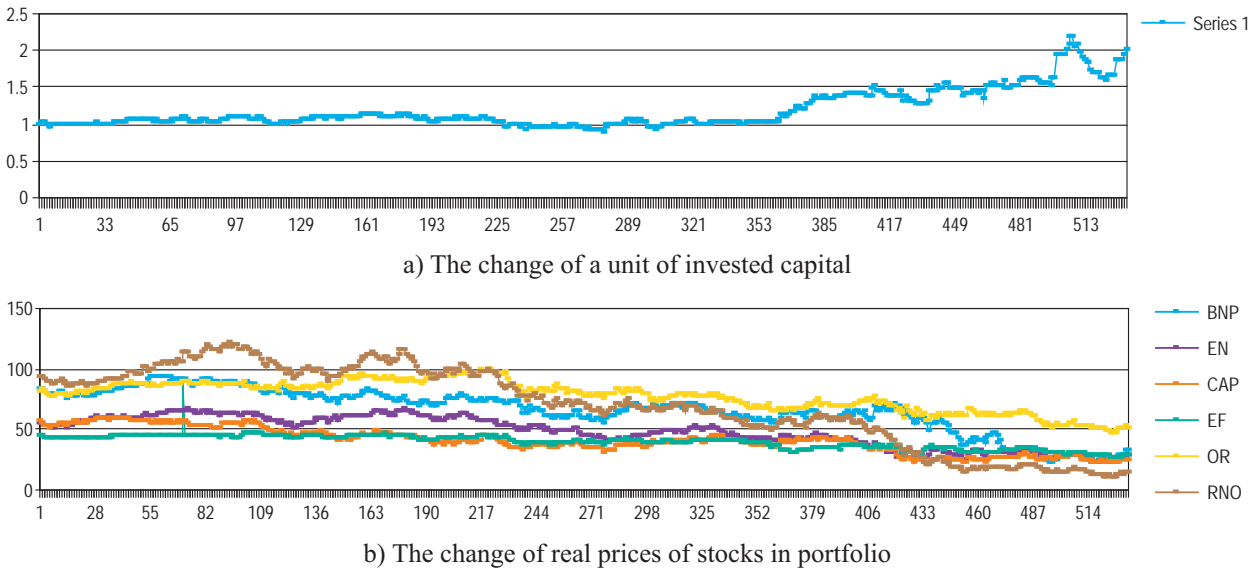


Fig. 7. The results of the experiment in France




Fig. 8. The results of the experiment in USA (NYSE)

5. Concluding remarks

In this paper we developed an artificial stock market model based on the interaction of heterogeneous agents whose forward-looking behaviour is driven by the reinforcement-learning algorithm combined with some evolutionary selection mechanism and economic reasoning. Other notable features of the model include knowledge dissemination and agents' competition for survival, detailed modelling of the trading process, explicit formation of dividend expectations and estimates of fundamental value, computation of individual reservation prices and best order prices, etc. At this stage of development, the model should largely be seen as a thought experiment that proposes to study financial market processes in the light of complex interaction

of artificial agents that are designed to act in an economically appealing way. Bearing in mind the uncertain nature of the model environment, mostly brought about by this same interaction, strategies followed by artificial agents seem to exhibit a good balance of economic rationale and optimisation attempts. Quite a strong emphasis on the model's economic content distinguishes this model from some other ASM models, which are most often based on evolutionary selection procedures and are sometimes criticised for lack of economic fundament.

Preliminary simulation results suggest that the market price of the stock in this model broadly reflects fundamentals but over- or under-valuation runs are sustained for prolonged periods. Both individual adaptive behav-

ior and the population level adaptation (evolutionary selection in particular) are essential for ensuring any efficiency of the market. However, market self-regulation ability is found to be weak. The institutional setting alone, such as the centralised exchange based on the double auction trading, cannot ensure effective market functioning. Even in the case of active adaptive learning, the market does not correct itself from erroneously perceived fundamentals if they are in the vicinity of actual fundamentals, which underscores the importance of market participants' beliefs in the market price dynamics. We also find a positive relationship between stock returns and changes in liquidity – there are indications that exogenous shocks to investors' cash holdings lead to strong changes in the market price of the stock.

Overall, this line of research seems promising. In our ongoing research, we aim at developing a version of the model suitable for calibration to empirical data. This requires simplification of some processes in the model, taking steps to ensure more effective and robust learning, etc. The noteworthy implication of the proposed study is that similar modelling principles could be expanded and applied for modelling of other markets, such as markets for goods or labour. More generally, intelligent adaptive agents could form the basis of applied dynamic macroeconomic models. It is very likely that in the future they will stand on the equal footing with the representative agent of dynamic general equilibrium models.

Parallel decisions search system, which is presented in the paper, exploits only part of the market – a certain amount of stocks. The system is based on assumption about market behaviour cognition, but the main model of market behaviour is admitted to be a multi-dimensional stochastic process, the identity of which regarding particular market is achieved with the help of stock prices, market indices and macro-economic data. Application of expert system allows to state that even under global financial crisis conditions distinct investment strategies are available, which guarantee long-term capital growth rates much higher than the general growth of the market.

References

Arthur, W. B.; Holland, J.; LeBaron, B.; Tayler, P. 1997. Asset pricing under endogenous expectations in an artificial stock market, in Arthur, W. B.; Durlauf, S.; Lane, D. (Eds.). *The Economy as an Evolving Complex System II*. Addison-Wesley, Reading, MA, 15–44.

Beltratti, A.; Margarita, S. 1992. Evolution of trading strategies among heterogeneous artificial economic agents, in

Meyer, J. A.; Roitblat, H. L.; Wilson, S. W. (Eds.). *From Animals to Animats 2*. MIT Press, Cambridge MA.

Kim, G.; Markowitz, H. 1989. Investment rules, margin, and market volatility, *Journal of Portfolio Management* 16: 45–52. doi:10.3905/jpm.1989.409233

LeBaron, B. 2000. *Evolution and Time Horizons in an Agent-Based Stock Market*. Available from Internet: <http://papers.ssrn.com/sol3/papers.cfm?abstract_id=218309>.

LeBaron, B. 2006. Agent-based computational finance, in Tesfatsion, L.; Judd, K. L. (Eds.). *Handbook of Computational Economics 2*: 1189–1233. Elsevier, Amsterdam.

Lettau, M. 1997. Explaining the facts with adaptive agents: the case of mutual fund flows, *Journal of Economic Dynamics and Control* 21: 1117–1148. doi:10.1016/S0165-1889(97)00046-8

Lux, T. 1995. Herd behaviour, bubbles and crashes, *Economic Journal* 105: 881–896. doi:10.2307/2235156

Ramanauskas, T. 2009. *Agent-Based Financial Modelling: A Promising Alternative to the Standard Representative-Agent Approach*. Bank of Lithuania Working Paper No. 3. Vilnius.

Rutkauskas, A. V. 2008. On the sustainability of regional competitiveness development considering risk, *Technological and Economic Development of Economy* 14(1): 89–99. doi:10.3846/2029-0187.2008.14.89-99

Rutkauskas, A. V.; Miečinskienė, A.; Stasytytė, V. 2008. Investment decisions modelling along sustainable development concept on financial markets, *Technological and Economic Development of Economy* 14(3): 417–427. doi:10.3846/1392-8619.2008.14.417-427

Rutkauskas, A. V. 2005. The double-trump decision management model in global exchange, *Ekonomika* 72: 84–104.

Rutkauskas, A. V. 2006. Adekvačiojo investavimo portfelio anatomija ir sprendimai panaudojant imitacines technologijas, *Ekonomika* 75: 52–76.

Rutkauskas, A. V.; Stasytytė, V. 2006. The double trump portfolio as the core of sustainable decision making strategy in currency markets, in *The 10th World Multi-Conference on Systemics, Cybernetics and Informatics*. Orlando, Florida, July 16–19, 57–62.

Samanidou, E.; Zschischang, E.; Stauffer, D.; Lux, T. 2007. Agent-based models of financial markets, *Reports on Progress in Physics* 70: 409–450. doi:10.1088/0034-4885/70/3/R03

Sethi, R.; Franke, R. 1995. Behavioural heterogeneity under evolutionary pressure: macroeconomic implications of costly optimization, *Economic Journal* 105: 583–600. doi:10.2307/2235019

Tay, N. S. P.; Linn, S. C. 2001. Fuzzy inductive reasoning, expectation formation and the behavior of security prices, *Journal of Economic Dynamics and Control* 25: 321–362. doi:10.1016/S0165-1889(00)00029-4

Watkins, C. J. C. H. 1989. *Learning from Delayed Rewards*. PhD Thesis, King's College.