Agent-Based Economic Modeling With Finite State Machines

Ivan Gnilomedov
St. Petersburg State University of Information Technologies, Mechanics and Optics
St.Petersburg, Kronverksky Prosp., 49

Sergey Nikolenko
Steklov Mathematical Institute, St. Petersburg
St. Petersburg, nab. r. Fontanka, 27

Abstract

We offer a novel approach to agent-based economic modeling. Previous work has modeled learning agents as neural networks, sets of fuzzy rules and other learning algorithms. In this paper, we present an approach based on representing an agent's production cycle as a finite state machine. We show that the finite state machine model offers natural representations for basic economic features such as complementary and substitute commodities and various production strategies with transition costs. Our experimental results show that the model behaves in perfect accordance with basic economic laws that shows the model's validity.

1. Introduction

Agent-based modeling has played an increasingly important role in understanding market behaviour. Existing research has been primarily centered on modeling financial markets with agents either following predefined strategies or learning the optimal strategy via different learning algorithms.

We aim to create a general equilibrium model representing a market of companies producing different commodities with different manufacturing processes, a model that would accurately reflect the behaviour of supply, demand, and prices in a real economy. A classical model of this market is the well-known Walrasian model for N different commodities produced with M different resources (production factors), where general equilibrium is achieved by maximizing each agent's profit; see (Black, 1995) for a detailed exposition. The general equilibrium model leads to nice mathematical properties (Walras' Law and others). However, this model assumes each agent is perfectly rational and computationally unbounded and the general equilibrium is static, while the real world market tries to achieve equilibrium in a constantly changing world. Thus, the need arises to offer agent-based models with limited and learning agents.

(Zimmermann, Neuneier, and Grothmann, 2001) offer an agent-based model of the *FX-Market* where agents base their decisions on incomplete information coming from a limited number of error-prone sources of information. The agents are modeled as error-correcting neural networks. (Kooths, Mitze, and Ringhut, 2004) and (Kooths, 1999) create a macroeconomic model where agents are modeled as neural networks with additional fuzzy rules representing knowledge about the economy. Other approaches have also been tested; for a detailed survey of agent-based economy models see (LeBaron, 2006), (Tesfatsion, 2006) and references therein.

Agent-based models with agents represented as finite state machines (FSMs) are widely used in computer animation (see, for example, (Rudomin, Millan, and Hernandez, 2005)), while learning FSMs has been already applied to automata-based programming developed by (Shalyto and Tukkel, 2001).

Developing systems of reactive agents with finite automata was suggested in (Naumov, Shalyto, 2003). The authors assert that their approach can allow to create systems of self-learning adaptive agents. The approach was further developed in (Shalyto, Naumov, Korneev, 2005), where interaction between objects was considered as interaction between finite automata. Finally, the theory was applied to creating a reactive multi-agent real-life environment, namely a system of robots who deliver items from one place to another (Yartsev, Korneev, Shalyto, Kotov, 2005). All robots in this system were controlled by finite automata logic.

However, as far as we know this is the first agent-based economy simulation where agents would be represented and trained as FSMs. In this paper, we fill the gap by constructing an economic model with adaptive agents based on finite state machines. The underlying FSM of an agent represents its production cycle. We have implemented a model economy with these agents, and results of our experiments were in good accordance with basic economic laws.

The paper is organized as follows. Section 2 describes the model itself: in 2.1, we show how to model a company with finite state machines, and 2.2 describes the market model in our approach and shows a concrete example of a production cycle. Section 3 lists the results of the experiments; we show that the model behaves just as the economic laws predict, thus establishing that the model is valid.

2. Description of the model

2.1. Modeling production with finite state machines

A *finite state machine* is a directed graph with captions on edges $FA = \langle V, E, \Sigma \rangle$ where V is the set of vertices, $\sum$ is the FSM's *alphabet* (the set of input events), and $E = \left\{ e = \langle v_1, v_2, \sigma \rangle \mid v_1, v_2 \in V, \sigma \in \Sigma \right\}$ is the set of edges (possible *transitions*). An edge $e = \langle v_{1,} v_{2,} \sigma \rangle$ means that transition from v_1 to v_2 with input event σ is possible. A finite state machine has two kinds of special vertices: a unique $s \in V$ is labeled as the *initial vertex* and some $\{ t_{0,} \ldots, t_K \} \subset V$ represent the set of *terminal* vertices. Sometimes it is reasonable to associate both income and outcome events with each edge. Then $Tr = \langle V, E, \Sigma, \Pi \rangle$,

$$E = \left\{ e = \langle v_{1,} v_{2,} \sigma, \pi \rangle \mid v_{1,} v_2 \in V, \sigma \in \Sigma, \pi \in \Pi \right\}$$

(these FSMs are usually called *transducers*). A *path* in a FSM is an ordered set

$$P = \langle v_{0,} e_{0,} v_{1,} e_{1,} \ldots, v_L \rangle \mid e_0 = \langle v_{0,} v_{1,} \sigma_{0,} \pi_0 \rangle,$$

$$e_1 = \langle v_{1,} v_{2,} \sigma_{1,} \pi_1 \rangle, \ldots, e_{L-1} = \langle v_{0,} v_{1,} \sigma_{L-1}, \pi_{L-1} \rangle$$

We call the word $ILabel(P) = \sigma_{0,} \sigma_{1,} \ldots, \sigma_{L-1}$ the *input label* of the path *P*, and by the *output label* of the path *P* we mean the word $OLabel(P) = \pi_{0,} \pi_{1,} \ldots, \pi_{L-1}$. A *cycle* is a path with identical start and finish: $v_0 = v_L$. We describe a complete working example of a production cycle FSM below; for more information on FSMs we refer to (Lothaire 2005).

We model the production cycle of each agent as a finite state machine. The state of the FSM corresponds to a certain stage of the production process. The initial state corresponds to the zero-stage of the production; the product then travels from one production department to another, as the FSM travels from one state to another.

Each state transition corresponds to performing a certain production stage and, as such, requires a certain amount of resources. We add an internal resource pool for each agent and consider the resources necessary to complete a certain production stage as the input event for the corresponding transition. The transition may take place only if enough resources are available.

A product travels from one production stage to another until it is ready. A complete product corresponds to a terminal state of the FSM. Then the product goes to a "warehouse" where it awaits deployment to the market. In the FSM terms, we model a certain production stage as a transition from a terminal state to the initial state; to this transition we associate the resources necessary to store and transport the complete product to the market.

Each production stage produces something; thus, we add to each edge a set of *output* resources that get produced during this transition. It is natural for finite state machines to have both input and output (entry and exit) actions; in our model, the entry action consumes resources, while the exit action produces new resources.

Besides a natural representation of a real-life company, the FSM formalism has several formal advantages. First, it easily allows an agent to be flexible. In the real world, an agent can reorganize his or her company to produce different commodities (at the very least, different kinds of a commodity). In the FSM model, we allow the FSM to be non-linear (to have forks). Depending on the path an agent takes, different commodities will be consumed and produced.

Second, while in the real world agents may adapt their strategies to changing market conditions, it takes time and resources to complete the adaptation. In the FSM model, this "transition cost" is modeled naturally: if a FSM is currently in a non-terminal state, it will take time and resources to reach the terminal state before the agent can choose a different path starting from the initial state.

Finally, in the real world the agents' behaviour is only suboptimal: agents do not possess complete information about the market and sometimes cannot compute the absolute best strategy. With this in mind, we model the agents as having rather simple strategies. An agent is modeled as a pair $\langle FSM, S \rangle$ consisting of a FSM and a strategy *S*.

2.2. Modeling the market

In this subsection, we describe how the market itself and interaction between agents are represented in our model. We model the market as a virtual "bulletin board" where each agent may post an offer indicating that he is willing to sell a certain amount of a certain product for a certain price. When an agent needs to buy a certain product, he queries the market for the selling offers on this particular product. To model insufficient information, we have the market to return only a certain random subset of the offers (the offers that this agent "knows of"). The agent then reviews these offers, chooses the most profitable, and satisfies them (buys the necessary product).

Note how the model naturally represents complementary goods and substitutes. In economics, commodities are called *complementary* if the cross-elasticity of their demand

$$E_c = \frac{Q_1^B - Q_0^B}{Q_1^B + Q_0^B} \frac{P_1^A + P_0^A}{P_1^A - P_0^A}$$

is positive, and the commodities are *substitutes* if the cross-elasticity is negative (here Q_0^B, Q_1^B – demand values for good *B* before and after the price of commodity *A* changed, P_1^A, P_0^A – prices on the commodity *A*). In the FSM model, complementary goods will often appear on consecutive edges of the production finite state machine, while substitute goods will appear on parallel edges or edges of parallel paths of the FSM. This results in that whenever an agent buys a certain commodity, his demand for the complementaries increases (he is halfway through the production, so he needs to complete it), while his demand for the substitutes decreases (he has already come through a certain path, and he will not return to a parallel edge in this cycle).

Finally, let us present a complete example of a production cycle in our model. The FSM belonging to an agent *D* is depicted on Fig. 1.

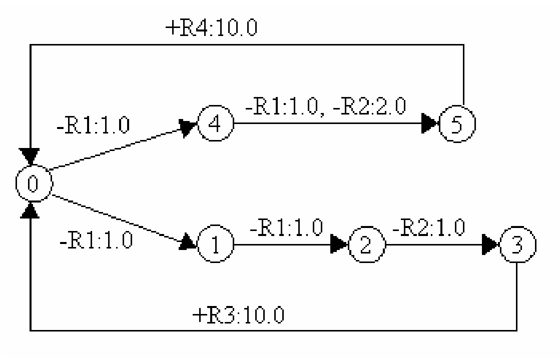

Fig 1. A sample production FSM

The initial state is marked as 0, while 3 and 5 are terminal states. Labels with – and + denote resource costs and resource outcome, respectively; the example has four different resources, from `R1` to `R4`. Here is a sample production cycle.

1. *Registering offers*. Suppose that other agents F_1 ... F_N have registered the following offers:

```
<R1, price=1.0, amount=12.0>,
<R1, price=2.0, amount=34.0>,
<R1, price=3.0, amount=45.0>,
<R2, price=10.0, amount=67.0>,
<R3, price=4.0, amount=89.0>.
```

2. *Deciding upon the production path*. Agent *D* analyses situation on the market and decides that the path {0, 1, 2, 3} is expected to be better than the path {0, 4, 5}. We do not specify here how agent *D* chooses his path; he may be guided by a reinforcement learning algorithm or by a complete statistical analysis of the market history.

3. *Buying resources*. To begin his production cycle, *A* needs to buy 1.0 of the resource `R1`. He sends a request for purchasing `R1`. The market randomly selects a subset of registered propositions, say,

```
<R1, price=1.0, amount=12.0>,
<R1, price=3.0, amount=45.0>,
```

and shows this set to *D*. Agent *D* chooses the most profitable proposition, with the price of 1.0, and makes his purchase.

4. *Transition*. Agent *D* now has the necessary resources to complete the first step of his strategy and move to state 1.

Note that after these steps agent *D* is completely committed to his chosen production path; he cannot switch to producing `R4` before completing this round of `R3`, even if in state 2 he suddenly discovers that the market is desperately lacking `R4`, and the prices could be exorbitant. We now skip a few steps as *D* gathers all necessary resources and completes his production cycle. Suppose that *A* has just made the transition from state 3 back to state 0 and obtained 10.0 units of `R3`.

5. *Registering a selling offer*. As *D* now has a surplus of `R3`, he will be willing to register an offer to sell `R3`. Again, we do not precisely specify the decision rule, but as input *D* should consider both his production cost and the market situation which he is able to peek by viewing the (random subset of) offers on `R3`.

After this step, the production cycle repeats. *D* is now able to choose another path.

3. Experiments

We have implemented the FSM market model and performed experiments that were aimed to check the basic economic facts about the model; they should serve as “sanity checks” for our modeling method. As a general result, the model turned out to be extremely viable, in clear-cut cases always performing just as the economic theory predicts. Thus, we expect it to have some predictive power as well, but this should be supported by further practical experiments. In this section, we describe the experiments we have carried out.

3.1. Production-possibility frontier

The production-possibility frontier is a curve that shows the maximal volume of producing a certain commodity dependent on the level of production of another commodity. This function should be decreasing, as increasing the level of production of a certain commodity should lower the amount of resources the market is ready to assign to producing other commodities. Besides, it should be convex, as in shifting from commodity *A* to commodity *B* the market would first shift the resources that are most useful for *B* and least useful for *A*.

Our experimental scenario included two commodities: *A* (automobiles) and *B* (blankets). Each agent is able to produce both commodities, and agent's FSM has several paths corresponding to different output ratio (automobiles)/(blankets) of the commodities that this agent produces during a production cycle. Moreover, each agent has a personalized production FSM that has its own (randomly varied) production costs and *A*/*B* ratios. A certain “command center” orders some agents to switch production from *A* to *B*, and the agents (with some delay, as described above) switch production; the remaining agents act in their own interest. Fig. 2 shows the results of this experiment: the production-possibility frontier is almost convex, with small fluctuations that can be related to suboptimal behaviour of the agents.

3.2. Overstocking

In this experiment we generated a world where each production cycle of each agent generates some surplus; in other words, each agents generates more resources than it consumes. In the real world, after an initial surge the overproducing warehouses would be full, and prices would experience a steady monotone decrease as the overproduction continues. As prices drop, some agents should become unprofitable and should quit the market (switch to the strategy of producing and consuming nothing), which gradually compensates overproduction and stabilizes the prices.

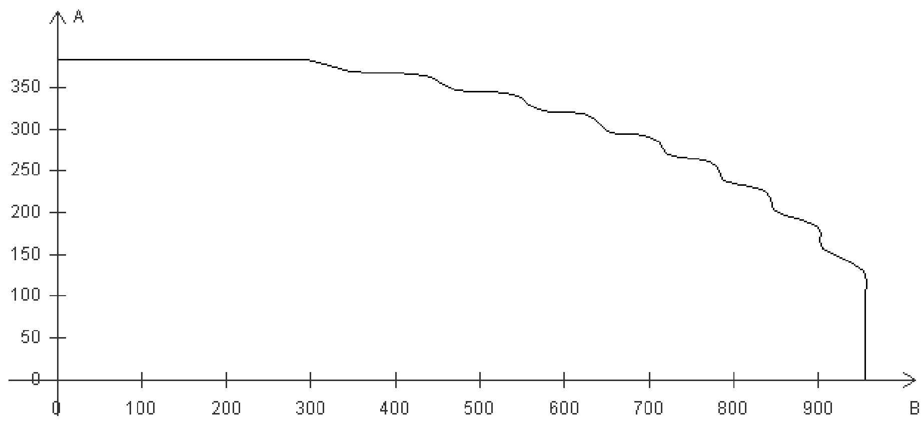

Fig 2. The production-possibility frontier

Fig. 3 depicts the results of the experiment; the graph on the left shows how prices behave with time, the graph in the middle shows how many agents are active (produce something) at this time, and the graph on the right shows the gross sales volume. Both curves behave just as predicted reach an equilibrium.

3.3. Equilibrium

In the real world, overproduction is natural, and it is naturally compensated by consumers who buy end products. To model this situation, we introduced a special agent (representing the consumers) that buys surplus resources. Fig. 4 shows that in this case, equilibrium is reached faster, and both sales and prices stabilize at higher levels.

3.4. Raw materials and end products

In this experiment we divided nine commodities of the model into two groups: {r0, r1, r2} and {r3, ..., r8}. We assume that production of a commodity depends strongly on other commodities of its group and to a much lesser extent on the commodities of the other group. At time 1000, we introduce an agent who buys lots of r0, thus raising prices. As a result, prices for r1 and r2 also rise, while prices for r3,...,r8 stay virtually the same. Fig. 5a shows the price graph.

3.5. Complementaries and substitutes

In the first part of this experiment, `r0, r1`, and `r2` are complementaries, that is, we increase the probability that they appear on consecutive edges of the production FSMs. An increase in the price of `r0` should cause a decrease in demand on `r1` and `r2`.

Fig. 5b shows what happens if at time 1000 we introduce an agent who buys lots of `r0`. The market reacts by rising all prices except for `r1` and r2: since producers buy less `r0`, they need less `r1` and `r2`, too. Fig. 6a shows the same experiment with only these three commodities left on the market, and Fig. 6b shows the results of an experiment with three commodities, among which `r0` and `r2` are complementaries: an increase in the price of `r0` causes an increase in the price of `r2`.

In the second part of this experiment, we introduce `r0, r1,` and `r2` as substitutes, that is, increase the probability that they appear on parallel paths in production FSMs. At time 600, we introduced an agent who buys lots of `r0`, thus increasing its price; as a result, agents begin to use less `r0` and more `r1` and `r2`. Fig. 6a shows the prices volume graphs for the three resources in this experiment. Fig. 7 shows the demand volumes for experiment with substitute commodities.

4. Conclusion

In this paper, we introduced a novel approach to modeling agents that act as producers and consumers on

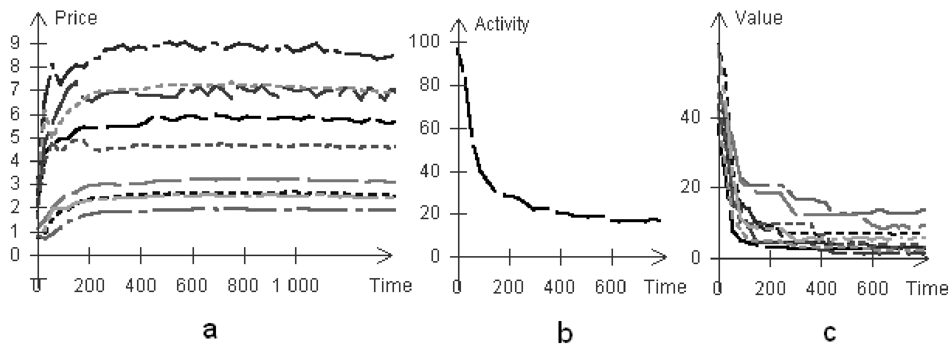

Fig 3. Overstocking: a – prices; b – active agents; c – gross sales

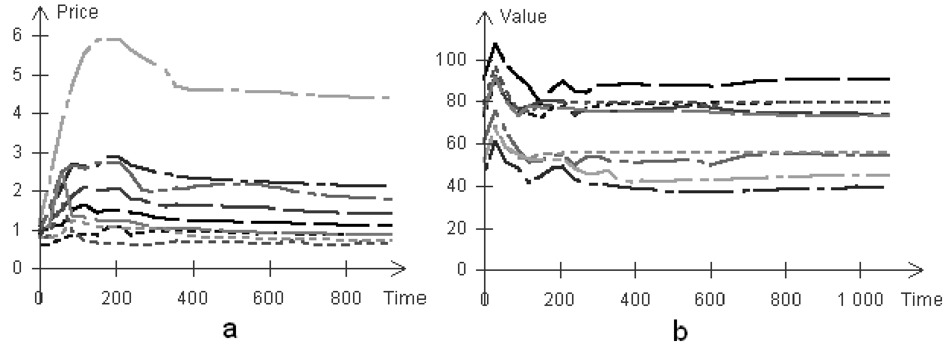

Fig. 4. Equilibrium: a – gross sales volume; b – prices

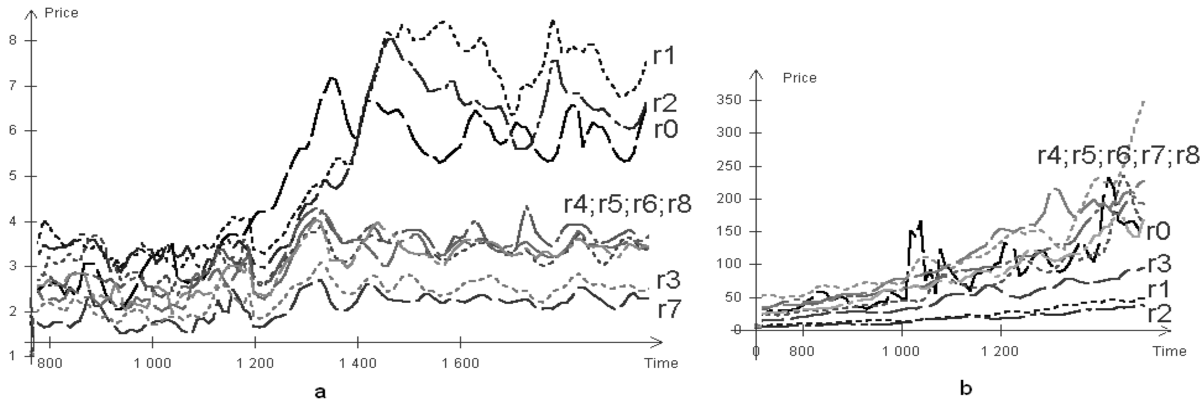

Fig. 5. Prices on a group of commodities: a – raw materials and end products, b – complementaries

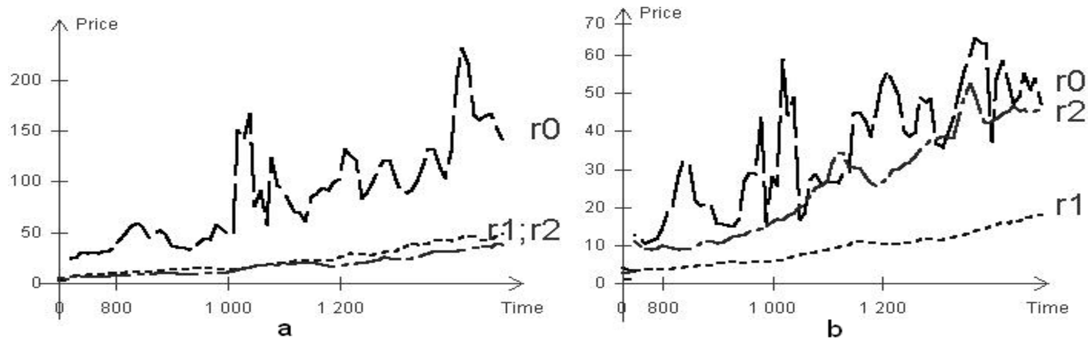

Fig. 6. Prices on complementary commodities: a – a group of three commodities; b – a group of two

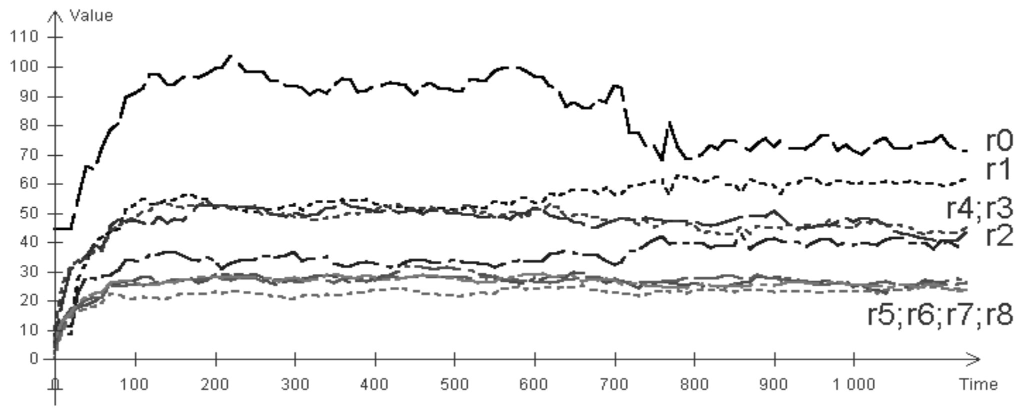

Fig. 7. Demand on substitute commodities

a market. We model an agent's production as a finite state machine with spent and produced resources as entry and exit actions. We have shown that a finite state machine is a natural way to model a company's production, and have shown how to model the basic properties of an economic agent and basic types of market commodities via finite state machines. We have implemented the proposed model and performed experiments in order to verify that the model works, and it works indeed: in all experiments the model behaved as expected.

As with any other model, we can point out the limitations of our approach. The primary limitation is that to create adequate results, the FSM agent-based model requires a large number of independent agents. This means that the model is applicable only to perfect competition markets and monopolistic competition markets.

This paper is, to a large extent, a proof of concept, evidence in support of the validity of the FSM model. Further work should deal with different learning algorithms and different pricing strategies for the agents. We plan to experiment how different strategies compete with each other and how well different FSM learning algorithms perform in this model. The model provides a natural competitive environment for testing various FSM learning algorithms against each other. And, of course, the ultimate test for our model would be to learn on real data (for example, on the stock market data) and test its predictive power.

References

F. Black. 1995. Exploring General Equilibrium. The MIT Press, Cambridge, Massachusetts.

G. Zimmermann, R. Neuneier and R. Grothmann. 2001. Multi-agent FX-Market Modeling Based on Cognitive Systems. In Proceedings of the Artificial Neural Networks, ICANN 2001, Lecture Notes in Computer Science, Vol. 2130, pp. 767–774.

S. Kooths. 1999. Modelling Rule- and Experience-Based Expectations Using Neuro-Fuzzy Systems. In Computing in Economics and Finance – Fifth Conference of the Society for Computational Economics, Boston.

S. Kooths, T. Mitze and E. Ringhut. 2004. Forecasting the EMU Inflation Rate: Linear Econometric Versus Non-Linear Computational Models Using Genetic Neural Fuzzy Systems. Advances in Econometrics, Vol. 19, pp. 145–173.

B. LeBaron. 2006. Agent-based Computational Finance, in Handbook of Computational Economics, L. Tesfatsion and K. Judd (eds.). North-Holland, Amsterdam, pp. 1187–1232.

L. Tesfatsion. 2006. Agent-Based Computational Economics: A Constructive Approach to Economic Theory, in Handbook of Computational Economics, L. Tesfatsion and K. Judd (eds.). North-Holland, Amsterdam, pp. 831–880.

I. Rudomin, E. Millan and Hernandez, B. 2005. Fragment shaders for agent animation using finite state machines. Simulation Modelling Practice and Theory, Vol. 13, Issue 8, pp. 741–751.

A. Shalyto and N. Tukkel. 2001. SWITCH-Technology: An Automated Approach to Developing Software for Reactive Systems. Programming and Computer Software, Vol. 27, No. 5, pp. 260–276.

L. Naumov and A. Shalyto. Automata Theory for Multi-Agent Systems Implementation Integration of Knowledge Intensive Multi-Agents Systems (KIMAS'03), IEEE Boston Section, 2003, pp. 65–70 .

A. Shalyto, L. Naumov and G. Korneev. Methods of Object-Oriented Reactive Agents Implementation on the Basis of Finite Automata Integration of Knowledge Intensive Multi-Agents Systems (KIMAS'05), USA, MA: IEEE Boston Section, 2005, pp. 460–465.

B. Yartsev, G. Korneev, A. Shalyto and V. Kotov. Automata-Based Programming of the Reactive Multi-Agent Control System Integration of Knowledge Intensive Multi-Agents Systems (KIMAS'05), USA, MA: IEEE Boston Section, 2005, pp. 449–453.

M. Lothaire. 2005. Applied Combinatorics on Words. Cambridge University Press.