Доклади на Българската академия на науките
Comptes rendus de l'Académie bulgare des Sciences

Tome 66, No 5, 2013

SCIENCES ET INGENIERIE

Automatique et informatique

STOCHASTIC CALIBRATION OF A NEUROCOMPUTATIONAL MODEL FOR ECONOMIC DECISION-MAKING

George Mengov

(*Submitted by Corresponding Member M. Hadjiski on January 29, 2013*)

Abstract

Using neuroscientific methods to study human decision-making is a new and exciting research area. Today it is dominated by observational approaches driven by the fMRI and other brain-imaging technologies, but it already links empirical data with concepts and findings from established paradigms such as utility theory and prospect theory. Developed in parallel over many decades, mathematical neuroscience can now offer a next generation of theories that help better understand the cognitive decision processes in general and those involved in economic choices in particular. However, the methods and techniques needed to implement sophisticated neuroscientific models for studying observational or experimental data are still in their infancy. This paper presents one effort in that direction and shows how the mathematical theory of reflex conditioning by Grossberg and Schmajuk [1] can account for important aspects of people's decisions in a laboratory experiment. A stochastically calibrated neurocomputational model emulated the cognitive mechanism for choosing one economic option among four, and was successful for 70% of the subjects in a sample of 80 people. Against a baseline probability of 0.25, its prediction rates for 9% of the people exceeded 0.75, and ranged between 0.37 and 0.70 for approximately 60% of the people. These results were very robust both for the model's impressive achievements as well as for its lackluster performances. We hypothesize that it predicted well only those persons who took decisions intuitively (as understood by Kahneman [2,3]), and failed with those who made efforts to develop a decision strategy. We anticipate our study to be followed by further steps in discovering the exact cognitive mechanisms of economic choice.

Key words: stochastic optimization, simulated annealing, READ neural circuit, gated dipole, economic choice

Introduction. What makes people choose one among several attractive options is a tantalizing question with obvious practical implications. Attempts at answering it have been made over the centuries both by psychology and economics. The past couple of decades witnessed unabating interest in the issue, fostered by the advancement of brain-scanning technology and neuroeconomics. However, the currently fashionable approaches are largely observational or phenomenological at best. They could only benefit from a more theoretical input, and one is available from the multitude of neurocomputational models that have been around for some time. Yet, modelling is more convincing when followed by experimental work. When a theory meets relevant empirical data, unexpected and interesting complications may arise that could instigate new research developments and chart new territories. One such example comprises the topic of the present paper.

We present a computer-based experiment of economic psychology built around the mathematical theory of Pavlovian and Skinnerian conditioning [1] and its central element, the READ (REcurrent Associative gated Dipole) neural circuit. In this study it is assumed that at least part of the explanation of human economic choice can be given by that theory. Its tenets are general enough to cover not only humans but also other animal species, and may be related to Kahneman's cognitive System 1 for intuitive thinking [2,3].

The experiment puts the participant in a situation to choose one offer among four, in order to receive a fictitious good called *omnium bonum* ("good for everybody", in Latin). It had to be nonspecific to avoid mental associations with real goods and services that could skew the participant's motivation. No transaction costs were involved in abandoning one vendor for another.

Each of the four suppliers offered and delivered different number of units of *omnium bonum*, whereby the riskiest one was also the most rewarding. The delivered quantities were chosen such that all suppliers remained competitive and were likely to form a distinct image in the eyes of the participant throughout the entire procedure of 20 rounds. A software application recorded people's actual choices and self-assessed satisfaction or disappointment in a Lickert scale at the exact moments when those two reporting actions took place. The experiment was conducted in 2010 with 80 students from the Faculty (School) of Economics and Business Administration at St. Kliment Ohridski University of Sofia. More details about the design, setting, and procedure can be found in [4,5].

Neurocomputational model. As it is typical for Grossberg's embedded neural models, READ combines three basic elements: short-term memory (STM) identical with neuron activity; medium-term memory (MTM) of neurotransmitter gates and synapses; long-term memory (LTM) storing information patterns contained in a population of neurons, active during certain time intervals. Because the neuron is interesting for this study only with regard to its activity, henceforth we will use "neuron", "neural activity", "activation", and "neural signal" as synonyms.

The READ model is shown in Fig. 1. In it, black circles represent neurons $x_1, \ldots, x_8, x_A, \ldots, x_D$, i.e. the STM; two squares designate MTM transmitter gates y_1 and y_2; semicircles represent long-term memories $z_{7k}, \ldots, z_{8k}$, where $k = A, B, C, D$ are four mutually exclusive categories, here associated with four economic alternatives. STM, MTM and LTM are described by three differential equations, one form of which is this:

***STM*: neuron activation.**

$$\frac{dx_i}{dt} = -a_1 x_i + (a_2 - x_i) J_i^+ - (x_i + a_3) J_i^- . \quad (1)$$

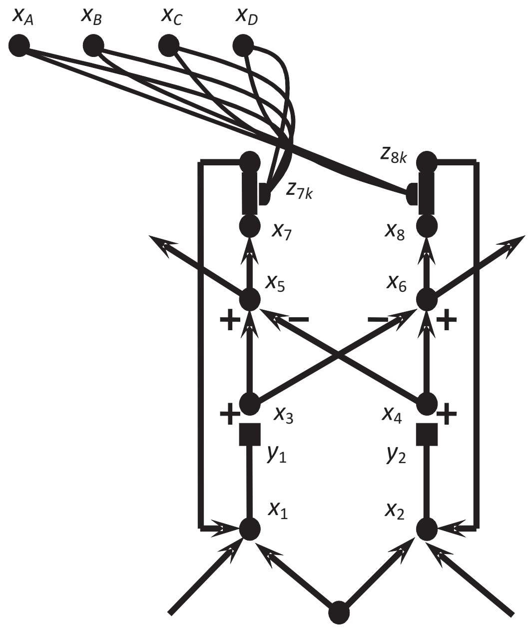

Fig. 1. READ model of psychological conditioning. The black circles are neurons sending out excitatory (+) and inhibitory (−) signals. The two squares designate gates emitting neurotransmitters. Structures with neurons abutted by semicircles represent biochemical long-term memories z_{7k}, z_{8k}, where $k = A, B, C, D$. Nodes x_k are neural representations of Suppliers A, B, C and D

Equation (1) is a variant of the classical Hodgkin–Huxley equation, cited here in its widely used dimensionless form. Term J_i^+ is the sum of all incoming signals from other neurons activating x_i. Similarly, J_i^- sums all inhibitory signals from other neurons. Constants a_1, a_2, a_3 are real and positive. All neuron equations used in this article's model are extensions or special cases of (1).

***MTM*: neurotransmitter gate.**

$$\frac{dy_i}{dt} = b_i(1 - y_i) - c_i x_i y_i . \quad (2)$$

Equation (2) describes neurotransmitter depletion and recovery in a neuron x_i, sending out a signal in response to received signals as per (1). Transmitter quantity y_i is a function of the strength of the emitted signal. In absence of external activation, the quantity y_i accumulates at fixed rate b_i until reaching its maximum of 1. The constants b_i, c_i are real and positive, and due to biological plausibility, in computer simulations may sometimes be considered slightly different for each gate i.

***LTM*: gated learning.**

$$\frac{dz_{ik}}{dt} = x_k(-h_1 z_{ik} + h_2 [x_j]^+). \quad (3)$$

Equation (3) describes how a single LTM component z_{ik} between neurons x_i and x_k changes under the influence of a third neuron x_j. Operator $[\cdot]^+$ denotes rectification: $[\xi]^+ = \max\{\xi, 0\}$. The cognitive mechanism can be explained with the help of Fig. 1 and is as follows. A signal x_k (with $k = A, B, C, D$) indicates which of the four suppliers is currently in the focus of attention. If, for example, this is Supplier A, then $x_A = 1$, and $x_k = 0$ for all the other k-s. That supplier's performance causes an emotion, represented by various degrees of satisfaction $[x_5]^+$ or disappointment $[x_6]^+$ in (3). In the role of x_i is neuron x_7 or x_8 (see Fig. 1). Synapse z_{ik} in (3) stores the emotional memory regarding how supplier k has satisfied or disappointed the customer. In particular, if the emotion was positive, z_{7k} is changed, otherwise z_{8k} is changed. The constants h_1, h_2 are real and positive. (In simulations, sometimes $h_2 = 1$.) Now, when the gating signal x_k is switched on, for example for Supplier A, and the customer is disappointed (positive x_6, zero x_5), neuron x_8 is active and memory component z_{8A} undergoes learning that reflects the input signals x_A and x_6. In other words, Supplier A is associated with disappointment this round.

Because neurons respond instantaneously to input signals, in computer models the STM equations are often solved at equilibrium. MTM and LTM are two or three degrees of magnitude slower, which requires analytical (preferably) or numerical solutions to their differential equations. In our case the former is possible for both (2) and (3). In particular, the exact analytical solution to (2), as integrated in our version of the READ computer model, is as follows:

$$(4) \qquad y_i(t) = y_i(t-1)\exp(-c_i x_i(t) - b_i) + \frac{b_i}{b_i + c_i x_i(t)}(1 - \exp(-c_i x_i(t) - b_i)),$$

where $i = 1, 2$. In (4), t is the discreet time with sampling rate sufficiently high to satisfy the Nyquist–Kotelnikov sampling theorem. Similarly, the solution to (3) is

$$(5) \qquad z_{ik}(t) = z_{ik}(t-1)\exp(-h_1 x_k(t)) + \frac{h_2}{h_1} o_l(t)[1 - \exp(-h_1 x_k(t))].$$

The quantity $o_l(t)$ is the READ-predicted consumer emotion and can be $o_1 = [x_5]^+$ or $o_2 = [x_6]^+$. Note that in (4) and (5), each of x_i, o_l and x_k is computed at moment t rather than $t-1$ because it is assumed to react instantaneously.

Stochastic calibration. A separate READ model had to be calibrated for each participant's empirical data, to explain and predict her or his economic choices. For the constants in the differential equations (1)–(3) such values had to be found that would make the model emulate the real decisions taken and satisfactions declared. Because our setting contained twenty rounds, we took the first twelve for calibration and used the remaining eight for testing.

By experimental design, the quantities of offered and delivered *omnium bonum* varied in each round so as to create a nonstationary process of economic growth, or growth followed by downturn. By adding this realistic element we subjected the READ model to a kind of stress-test, and expected it to cope well due to its biological origins. Yet, it is natural that predictions of the test sample could not be as good as in a hypothetical experiment with a stationary process.

The model was calibrated with simulated annealing – a method for stochastic optimization. In a previous study [6] we defined an objective function, optimized with respect to both emotional self-assessments and supplier choices, each carrying the same weight. This time we chose to account better for the process nonstationarity by maximizing an objective function J of the form

$$J = I_{F8} + 5I_{M4} + R\left(\mathbf{DS}(\mathbf{t}_{DS}^{(F12)}), \mathbf{o}(\mathbf{t}_{DS}^{(F12)})\right). \tag{6}$$

In equation (6), I_{F8} is the number of correct choices READ made in the first eight rounds ($F8$) of the calibration sample; I_{M4} is the number of its correct choices in rounds Nos 9–12, i.e. the last four in the calibration sample (the middle four, or $M4$, in the entire sequence of twenty). Multiplying I_{M4} by five substantially enhances the importance of the final calibration rounds, thus mitigating as much as possible the nonstationarity of the process. The weight coefficient was chosen to be 5 heuristically but not arbitrarily – it had to ensure a good balance between the two subsets in the calibration sample. Due to equation (6), for example a combination of all correct $M4$ choices and totally incorrect $F8$ choices ($I_{F8} + 5I_{M4} = 20$) was valued more than another combination of all correct F8 plus half-wrong M4 (sum equal to 18). In this line of reasoning, combinations of the same number of correct choices were treated differently depending on their $F8$:$M4$ configuration and, say, four correct choices might occasionally be preferred to ten, as in the above example. However, there is no conceptual problem here because the simulated annealing procedure is theoretically guaranteed to converge to the global maximum of J regardless of J's form. In addition, this strategy was an effective safeguard against falling in traps of local error minima.

The term $R\left(\mathbf{DS}(\mathbf{t}_{DS}^{(F12)}), \mathbf{o}(\mathbf{t}_{DS}^{(F12)})\right)$ is the correlation between vector $\mathbf{DS}$ of a participant's answers to the disappointment-satisfaction question, and vector $\mathbf{o}$ of their READ-computed forecasts, all at the recorded decision moments $\mathbf{t}_{DS}^{(F12)}$ during the first 12 rounds. Now, in the frequent situations with two identical $I_{F8} + 5I_{M4}$ sums, R tips the balance in favour of the solution better accounting for the participant's emotions.

Finally, let us explain what the factors forming the prediction criterion were. We based the READ decision-making mechanism on three factors: F_1: The momentary emotional reaction to the current *omnium bonum* offers; F_2: The experiences with all suppliers accumulated in the customer's long-term memories, i.e., the suppliers' reputations; F_3: the particular emotion remembered after the last dealing with a supplier regardless of how far in the past it happened.

All three factors were scaled to fit in the interval $[0,1]$. In more detail, F_1 is a normed nonlinear function of the neural activations x_7 and x_8, as they respond to each of the four offers. F_2 is the logistic function of the ratio z_{7k}/z_{8k}. Finally F_3, the last remembered emotion about a supplier, was rescaled from the interval $[-4,+4]$ to fit in $[0,1]$. The three factors were weighted by positive constants d_1, d_2, d_3 to form the following decision criterion:

$$D_k = d_1 F_1 + d_2 F_2 + d_3 F_3, \quad (7)$$

where $d_1 \in (0,1)$, $d_2 \in (0,1-d_1)$ and $d_3 = 1 - d_1 - d_2$. Equation (7), with its constraints on d_1, d_2, d_3 ensured that each factor could vary widely within the open interval $(0,1)$ at the expense of the other two, should the simulated annealing procedure require so. Criterion D_k was computed for each $k = A, B, C, D$ and the largest D_k was proclaimed winner, i.e., READ's choice of supplier in the current round. The constants d_1, d_2 were subjected to stochastic calibration alongside nine other constants from the differential equations.

Results and discussion. Each person's empirical data were fit to the READ model via the stochastic optimization procedure by computing typically 4000–20 000 solutions to the system of equations. This was done at least twice for each of all 80 people in the sample. The results can be considered optimal in the sense that it is virtually impossible to achieve higher values for the objective function.

In line with statistical learning theory [7], it was expected that more calibration effort will lead to consistently higher prediction accuracy in the calibration sample, and will be accompanied by a bow-like effect for the test sample as it is shown in Fig. 2. The present study can be put into a subcategory of the theory whereby the model complexity is fixed, but the number of times the equations are solved can vary.

An essential part of our empirical findings is summarized in Fig. 3. The y-axis in each plot shows the prediction error in a zero-to-one scale, with "zero" meaning that all 12 calibration choices on the solid line are correctly guessed, or more precisely, fit by READ, while 1.0 means not a single success at all. Of course, the prediction is genuine only for the test sample (dotted line), where zero error stands for full eight correct guesses in rounds Nos 13–20. Each dot on the solid line represents a new maximum of J as computed by (6). The x-axis shows the number of maxima achieved during a simulated annealing run, but not the actual number of computations. Although the plots show all maxima at equal distances, in practice the game of chance may separate every two neighbours by several to several thousand computations. The black squares along the dotted lines designate errors computed for the test sample each time the algorithm achieved a new maximum of J for the calibration sample. In this way it was possible to monitor how well the process conforms to statistical learning theory.

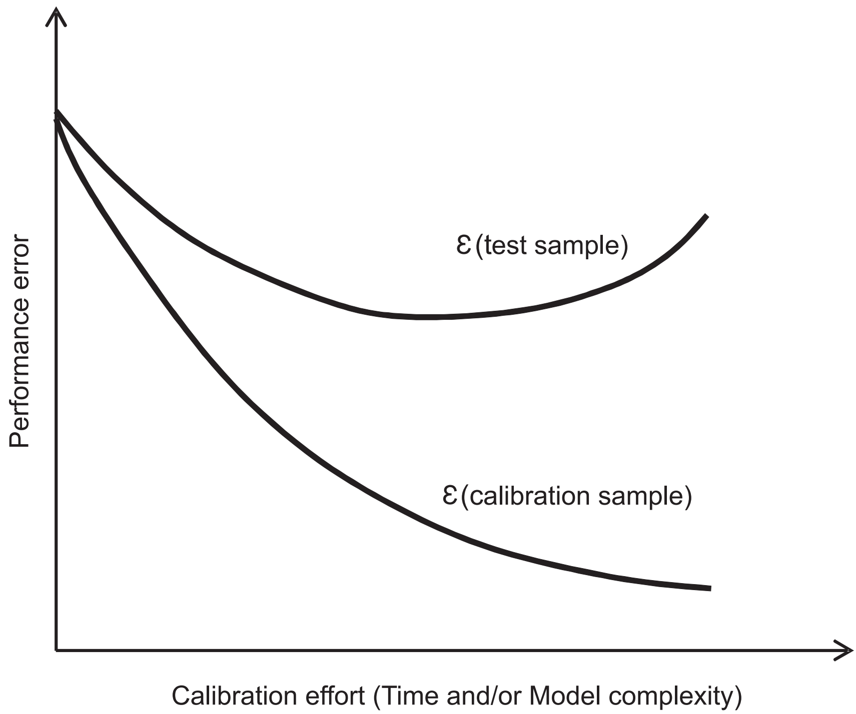

Fig. 2. Theoretical error curves in calibration and test samples

Figure 3*a* offers a good match between the computer simulation and theory. Occasionally, a new dot (better J) leads to an increase (instead of reduction) of the calibration error due to the objective function discontinuity explained in the previous section. Yet, the general picture is in line with the expectations. However, Fig. 3*b* looks less convincing, while Fig. 3*c* shows two cases of the unchanging test sample error no matter what the calibration success has been.

What conclusions can be drawn from this variegated picture? As the plots of Fig. 3 show, the ideas of statistical learning theory provide useful guidance, but should be accepted only as statistical tendency. Indeed, the overall success was 0.5490 for the calibration sample and 0.4224 for the test sample. (These numbers are averages on 80 people.)

In practice, the graphical patterns emerging from thousands of computations may depart substantially from what was expected (Fig. 2). Of course, the theory is not to blame – a number of additional issues complicate the situation. Among these may be listed: (1) the relatively small number of rounds in the experiment as well as the small samples for the experimental conditions; (2) the limited number of computations – thousands rather than millions, leading to (3) too steep a descent in the stochastic optimization procedure.

On the other hand, we obtained some very interesting results. Overall, in the test sample READ's prediction was useful for 70% of the people, while for

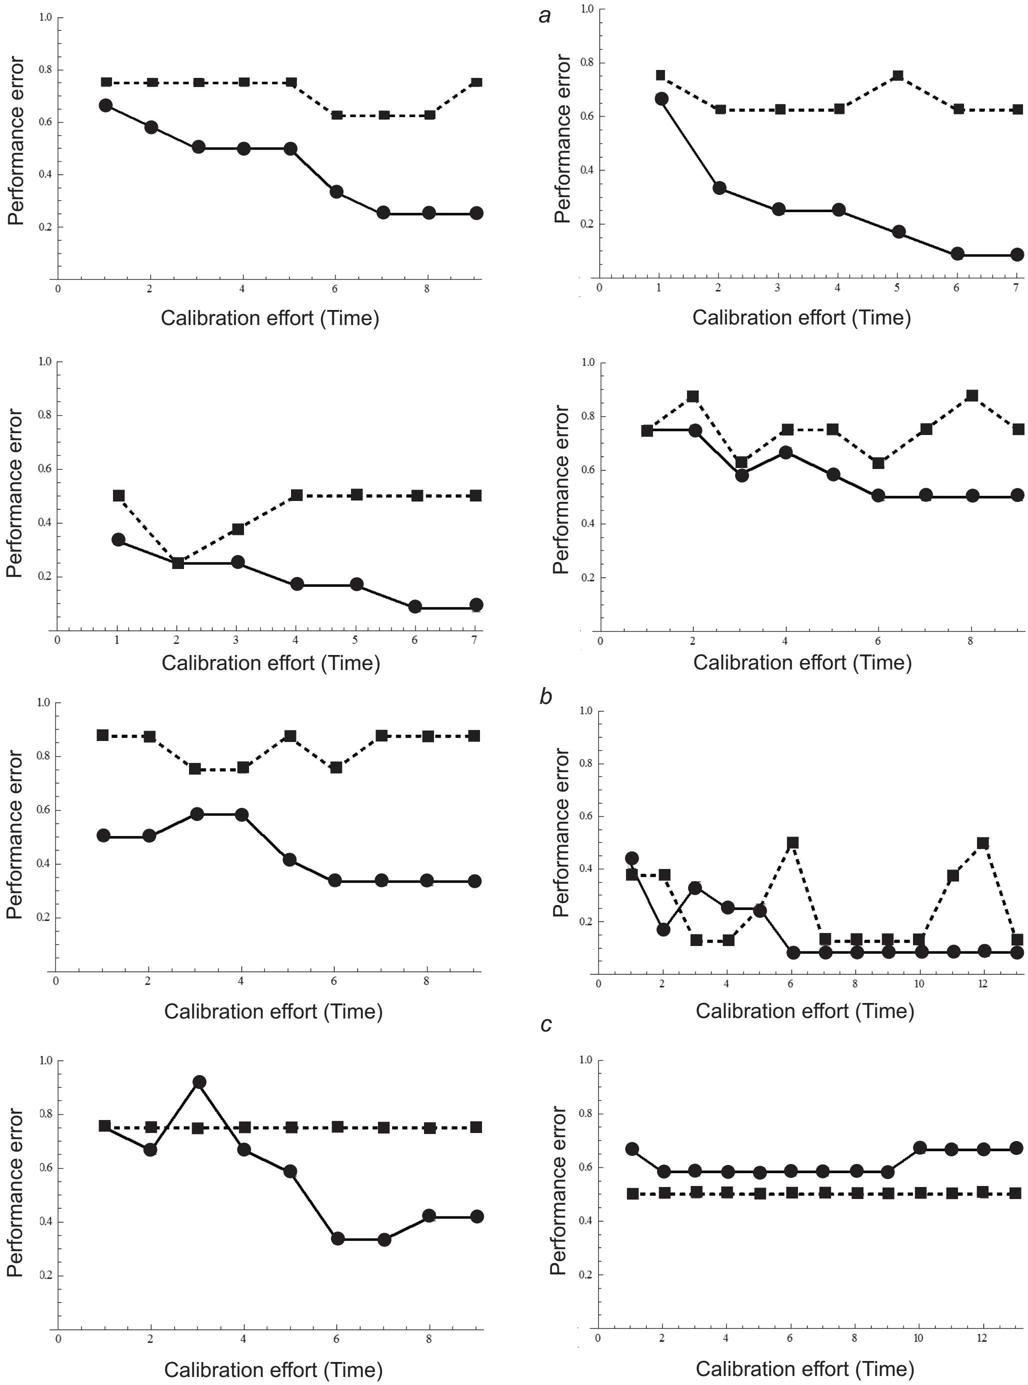

Fig. 3. Empirical error curves in calibration sample (solid line) and test sample (dotted line): (*a*) curves in agreement with Fig. 2; (*b*) curves less similar to the theoretically expected; (*c*) no improvement of choice prediction in the test sample

the remaining its correct guesses failed to rise above 25%. It achieved at least 50% successful predictions for 42.5% of the people, and rose to 75% or more for seven participants out of 80. For two individuals, of whom one is presented in Fig. 3*b* (right plot), READ provided very precise forecasts – around 90% – both in calibration and test samples. Remarkably, for them, just like for everybody else, the simulated annealing procedure consistently yielded the same final results.

The most interesting problem was that although the prediction success (number of correct guesses) was stable for each person across many runs with thousands of computations, the final values of the eleven constants varied substantially. Apparently they influenced and compensated each other much like in any other statistical learning system (e.g. neural networks of various types) with even a moderate number of parameters. However, for 70% of the people for whom the model was useful it was established that statistically, i.e. on average, the supplier reputations (Factor F_2) were twice more influential for the decisions than were the perception of current offers (F_1) and the last remembered emotion about a supplier (F_3). The latter two were about equally influential. That finding must be taken with caution though, because the three factors' weights had mean-to-standard-deviation ratios between 1.09 and 1.71, which indicates too much variance.

Conclusions. A general conclusion may be that READ performed within a wide range – from a huge success with about 9% of the people, to prediction levels of 37–70% for about 60% of the people, to no achievement at all with another 30%, with all results being very solid. Although this study cannot answer why that happened this way, we can hypothesize that READ predicted well only those people who took decisions intuitively and made no effort to develop a strategy. This suggestion naturally leads to the next research task – a repetition of the present experiment, complemented by psychological tests and questionnaires to find out what cognitive mechanisms and which personality traits influence decision-making both in the laboratory and at large.

Outside the computational issues, it may be that the human brain does not maintain the mix of relevant factors constant, but rearranges them ad hoc for each decision. If this be so, then models such as this one will be able to show only what might be happening and possibly discover some statistical laws, but will be unable to guarantee deterministic solutions.

REFERENCES

[1] Grossberg S., N. Schmajuk. Psychobiology, **15**, 1987, No 3, 195–240.

[2] Kahneman D. The American Economic Review, **93**, 2003, No 5, 1449–1475.

[3] Kahneman D. Thinking, Fast and Slow, New York, Farrar, Straus and Giroux, 2011.

[4] Georgiev N., G. Mengov. Annuaire Univ. Sofia Fac. Econ. Gest., **10**, 2012 (in print, in Bulgarian).
[5] Mengov G., N. Georgiev. Compt. rend. Acad. bulg. Sci., **66**, 2013, No 3, 415–422.
[6] Mengov G., H. Egbert, S. Pulov, K. Georgiev. Neural Networks, **21**, 2008, No 9, 1213–1219.
[7] Vapnik V. IEEE Trans. on Neural Networks, **10**, 1999, No 5, 988–999.

Department of Statistics and Econometrics,
Faculty of Economics and Business Administration
St. Kliment Ohridski University of Sofia
125, Tsarigradsko chaussée Blvd, Bl. 3
1113 Sofia, Bulgaria
e-mail: `g.mengov@feb.uni-sofia.bg`