Supporting Financial Decision Making by an Intelligent Agent Estimating Greed and Risk

Tibor Bosse[1], Ghazanfar F. Siddiqui[1,2], and Jan Treur[1]

[1]Vrije Universiteit Amsterdam, Department of Artificial Intelligence,
De Boelelaan 1081, 1081 HV Amsterdam, The Netherlands
http://www.few.vu.nl/~{tbosse, ghazanfa, treur}
{tbosse, ghazanfa, treur}@few.vu.nl
[2]Quaid-i-Azam University Islamabad, Department of Computer Science, 45320, Pakistan

Abstract — **In the area of financial decision making it is more and more acknowledged that psychological states and characteristics play an important role, for example feeling insecure in relation to financial risks, and being greedy in relation to opportunities to obtain serious gains. This paper presents an agent model of human decision making behaviour in economic situations, incorporating a human's greed state and personality characteristic concerning risk. The model provides a basis for the development of personalised intelligent agents that support a person in financial decisions. To evaluate the model a number of simulation experiments have been performed, which illustrate the model's ability to show behaviour of different types of personalities.**

I. Introduction

Financial decision making is not a trivial task for humans. In 1979, Kahneman and Tversky [6] claimed that humans do not behave rationally when they have to decide between alternatives that involve risk, as, for example, in financial situations. Also in other literature it has been argued that theories of economic decision making need to incorporate psychological factors such as greed and fear [4, 10, 12, 14]. The current paper is part of a project that aims to develop a personalised intelligent agent which supports persons in making investment decisions. The main goal of the project is to develop an agent that has insight in the individual psychological states and characteristics of persons using financial applications, and is able to exploit this insight in order to provide appropriate support, both in following these states and characteristics in a personalised manner, and in encouraging reflection by the person through mirroring his or her states and decision making processes. In order to develop such a support agent, as a first step, a solid computational model of human decision making in financial context is needed. The development of such a model is the main contribution of the current paper. The model takes the main principles underlying the Modern Portfolio Theory (MPT) [5, 13] as a basis, and extends these with mechanisms to incorporate psychological factors (inspired, among others, by [2, 7, 9, 10, 12]). In [3] it is shown how the model can be incorporated in virtual agents, and some human experiments are reported.

The remainder of this paper is structured as follows. In Section 2, a global overview of the agent model for economic decision making is presented, and in Section 3 this model is formalised. Some simulations of the model are discussed in Section 4, and Section 5 presents a mathematical analysis of equilibrium states of the model. Section 6 discusses how the risk profile is adapted to the person. Finally, Section 7 provides a summary of the work and a discussion about future research.

II. Overview of the Agent Model

An overview of the agent model for financial decision making and its interaction with the world is depicted in Figure 1. Here input states are depicted at the left hand side of a box, internal states within the box, and output states at the right hand side. The arrows indicate (causal) relationships between states. As shown in Figure 1, the agent model includes a notion of *greed* as a mental state of an individual. This greed is assumed to be a dynamic state, which is continuously influenced by two other concepts, namely the *observed results of investments* and the individual's *personality profile* concerning risk taking or risk avoidance. The former is considered to be the result of (dynamic) real world investments, depending, among others on the state of the world's economy, and risks taken in product choices in the past, whereas for the moment the latter is assumed to be a static characteristic of an individual. For example, persons who are more risk seeking than others may more easily get disappointing outcomes; if many investments have failed recently, persons are more likely to reduce their greed and as a consequence take less risk. This main idea is similar to the ideas in existing literature such as [2, 7, 9, 10, 12].

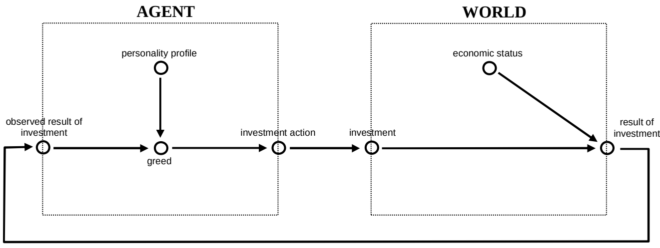

FIGURE 1. THE AGENT AND ITS INTERACTION WITH THE WORLD

To create an appropriate economic context, one particular type of investment decision is considered, namely the task to choose a financial product characterised by two factors concerning risk and expected gain. Thus the set of products is represented using a standard risk/return curve as proposed in the literature on Modern Portfolio Theory (see Section 3 for more details). After an *investment* decision has been made, the selected product is transferred to the world, where the *result of the investment* is determined, which depends not only on the selected product (in the sense that a more risky product has a lower probability to result in some return), but also on the *economic state* of the world (when the economy becomes more positive, then the probability of receiving some return increases). Note that, for simplicity, the current model considers the economic status as an external variable, although in reality this variable may depend on many other factors, such as the economic behaviour of other agents in the system (cf. [2]). The results of investments are observed by the individual, which completes the loop between agent and world.

III. FORMALISING THE AGENT MODEL

In this section, the global relations presented above are formalised in terms of mathematical specifications. In the model 10 products are used. Individuals are able to choose between these products by taking two characteristics into account, namely *risk* and *expected return*. To create a realistic range of products, the following parabolic equation is used for the relation between expected risk X and expected return Y (cf. [5, 13]): $X=aY^2+bY+c$ with $a=1$, $b=-0.1$ and $c=0.1$. The graph is shown in Figure 2.

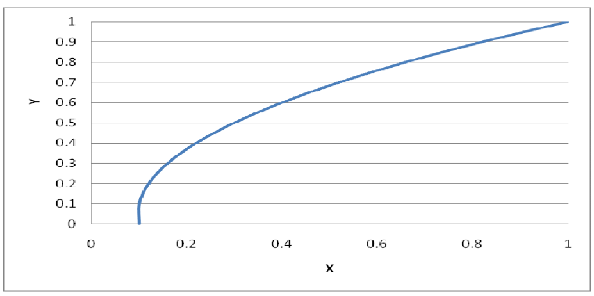

FIGURE 2. EXPECTED RETURN/RISK CURVE

The idea is that more greedy persons will select products that are further to the right hand side of the curve.

Updating greed. The dynamics of greed are modelled as follows:

$$G(t+\Delta t) = G(t) + \beta(((p+1)/2)*E - G(t))\Delta t$$

Here $G(t+\Delta t)$ is the updated greed, $G(t)$ is the old greed, E is the world event concerning the return on the earlier chosen product, and p is the individual static risk profile (0 means that the individual is low risk taking and 1 means the individual is high risk taking). The persistency factor β is the proportion of the old greed that is taken into account to determine the new greed. Simulation tests indicated that $\beta = 0.1$ gives realistic results. The values of G, E and p are in the range between 0 and 1. Thus, the underlying idea of the formula is that a person may show more greedy behaviour if the individual risk profile is more risk taking, and if more positive experiences were received in the recent past (see also [2, 7, 9, 10, 12]).

Selecting a product. A product is selected based on the greed and the personality characteristic p. As a first step the following factor r is determined:

$$r = (1/G)-1\ ) / (2*(p+1))$$

This r is taken as the required slope of the curve (depicted in Figure 2) for the product to be chosen, according to Modern Portfolio Theory. The choice of the product is made as follows. For each of the considered products (X, Y) the following is calculated:

$$Z(X, Y) = Y - r*X$$

Then the product (X, Y) is chosen with the highest $Z(X, Y)$. This product is the closest approximation of the point at the curve with slope r.

Determining the return. The model for the return E of the investment is as follows; here W is the state of the world (taken between 0 and 1), and (X, Y) is the selected product:

1. *Generate a random number C between 0 and 1 (both inclusive)*
2. *IF* $C \geq X * (1-W)$ *THEN* $E = Y$
3. *IF* $C < X * (1-W)$ *THEN* $E = 0$

This shows, for example, that when the state of the economy $W = 1$, there is no risk to have no return, and when $W = 0$ this risk has probability X.

IV. SIMULATION RESULTS

Based on the agent model introduced, a number of simulation experiments have been performed under different parameter settings of p. In Figures 3 to 5, time is on the horizontal axis and G, W, E and the average profit received by the person are represented on the vertical axis.

In all the simulation experiments, the initial value of greed is *0.5* and $\beta = 0.1$. For the value of *W*, a scenario has been established that is based on existing empirical data. For these data, the global Gross Domestic Product data over the period 1969-2008 have been taken from [17].

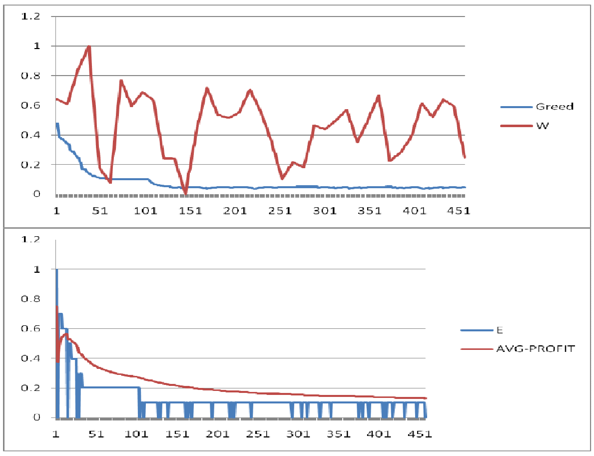

FIGURE 3. SIMULATION RESULTS FOR $p = 0$

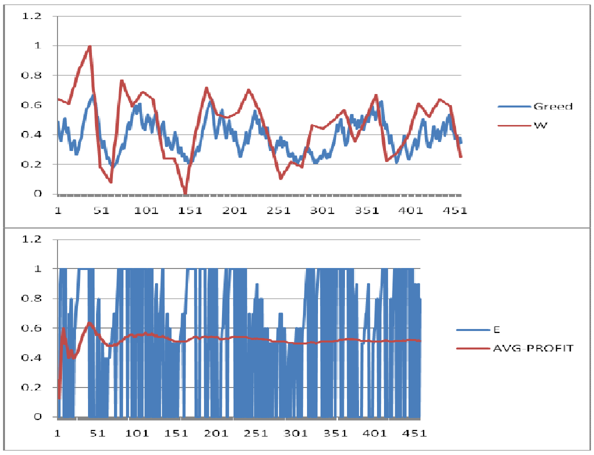

FIGURE 4. SIMULATION RESULTS FOR $p = 0.5$

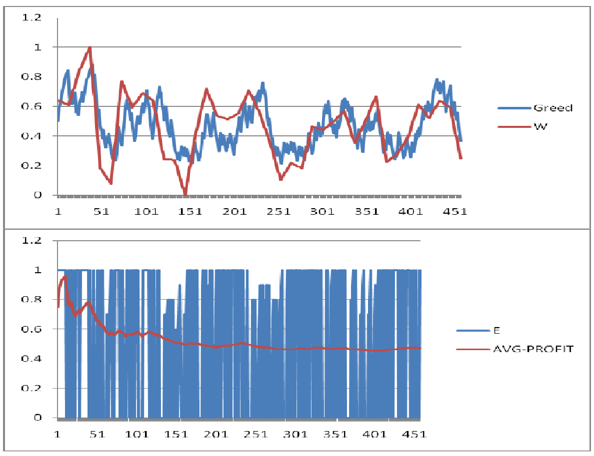

FIGURE 5. SIMULATION RESULTS FOR $p = 1$

As illustrated by these simulations, in most cases the person adapts its greed *G* to the status of the economy *W*. Thus, the persons quickly determine which level of greed is most appropriate in which situation. Only in the cases with very low *p*, after a while the greed becomes so low that the person cannot recover from that anymore (because it will never have very positive experiences). When comparing the different graphs with each other, the values of the events (and thus also the average profit) are highest in the case where $p=0.6$ (not shown due to space limitations). In cases with a higher *p*, the average profit slightly decreases, although not much.

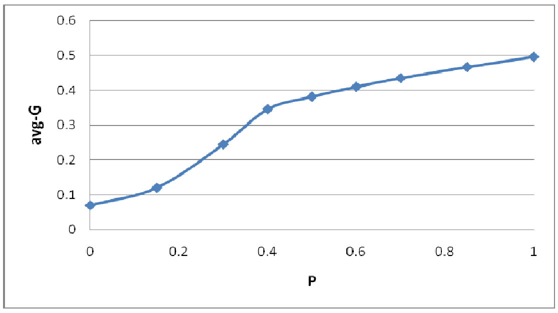

FIGURE 6. RELATING AVERAGE GREED AND *p*

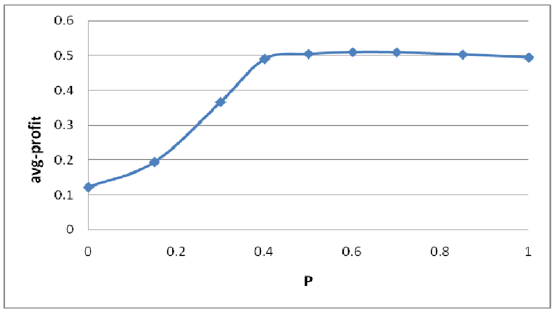

FIGURE 7 RELATIING AVERAGE PROFIT AND *p*

In addition, a larger number of simulation runs have been performed to see the effect of *p* on the average greed and average profit over the simulation. For this, 50 simulations were taken for each value of *p*, and graphs were plotted between *p*, average greed and average profit. In Figure 6 and Figure 7, *p* is on the horizontal axis and average greed or average profit is on the vertical axis. It is evident from Figure 6 that as *p* increases, the average greed also increases.

Furthermore, Figure 7 shows roughly that the average profit increases with an increase in the value of *p*. The average profit increases sharply from $p = 0$ to $p = 0.4$, and then becomes more or less stable. The point where average profit is maximal lies at $p = 0.6$. Hence, apparently this is a scenario where it is optimal to have a personal risk profile of *0.6*. In other scenarios (e.g., with a very bad economy), other risk profiles may be more beneficial.

V. Mathematical Analysis

For any given fixed (constant) values of p and W, the process approaches an equilibrium state, as was found by various example simulations. In this section a mathematical analysis of equilibrium states of the model is discussed. Knowledge of how these equilibrium states are, is also useful to analyse functioning of the process for changing W, as after each change in W, the process will aim to reach an equilibrium state for this new value.

The general setup is based on the continuous model given by the following differential equation for G:

$$dG/dt = \beta(((p+1)/2)E - G)$$

Each time the product is chosen on this curve for which

$$dY/dX = r \text{ or } dX/dY = 1/r$$

according to Modern Portfolio Theory, with

$$r=((1/G)-1)/(2(p+1)).$$

Relation between G and the product chosen From the product curve equation it follows that $dX/dY = 2aY + b$. Therefore:

$$\begin{aligned} 2aY + b &= 1/r \\ &= 2(p+1)/((1/G)-1) \\ &= 2(p+1)G/(1-G)= 2(p+1)G/(1-G). \end{aligned}$$

This can be used to express Y in G thus obtaining:

$$Y = ((20p+19)G + 1)/20(1-G)$$
$$X = Y^2 -0.1Y + 0.1$$

Some special cases are:

$$G = 0 \Rightarrow Y = 1/20$$
$$p = 0 \Rightarrow Y = (19G + 1)/20(1-G)$$

Determining an expected value for E An equilibrium for G has to satisfy $dG/dt = 0$. Given the differential equation this is equivalent to

$$((p+1)/2)E = G.$$

The events considered as outcomes E depend on a random number C in $[0, 1]$: if $C \geq X.(1-W)$ then $E=Y$, else $E=0$. Therefore $E = 0$ with probability $X.(1-W)$, and $E = Y$ with probability $1 - X.(1-W)$. From this the expectation value $ExpVal(E)$ for E is determined as:

$$ExpVal(E) = (1 - X.(1-W))\, Y$$

As E itself is always fluctuating due to the randomness, no real equilibrium can occur. Therefore, for an equilibrium analysis it is better to take the expectation value $ExpVal(E)$ instead of E itself.

The equilibrium equations An estimation for an equilibrium can be obtained when for E its expectation value is taken. Then the equation for $dG/dt = 0$ is: $(((p+1)/2)*ExpVal(E) = G$. This can be rewritten into $G = 0.5\ Y\ (p+1)(1 - X.(1-W))$. This equation can be combined with the previous one to form the following set of three equilibrium equations with three unknowns:

$$G = 0.5\ Y\ (p+1)(1 - X.(1-W))$$
$$X = Y^2 -0.1Y + 0.1$$
$$Y = ((20p+19)G + 1)/20(1-G)$$

An example solution for $p=1$ and $W=0$ is: $G = 0.226$, $Y = 0.265$, $X = 0.144$ (simulated). Some other cases are:

$$p = 0 \Rightarrow G = 0.5\ Y\ (1 - X.(1-W))$$
$$p = 1 \Rightarrow G = Y\ (1 - X.(1-W))$$
$$W = 1 \Rightarrow G = 0.5(p+1)Y$$
$$p = 0\ \&\ W = 1 \Rightarrow G = 0.5\ Y$$
$$p = 1\ \&\ W = 1 \Rightarrow G = Y$$

Overview of solutions of the equilibrium equations
A large number of simulations have been performed for different values of p and W. The outcome of the equilibrium values for G found is shown in Figures 8 (G against p for different values of W) and 9 (G against W for different values of p).

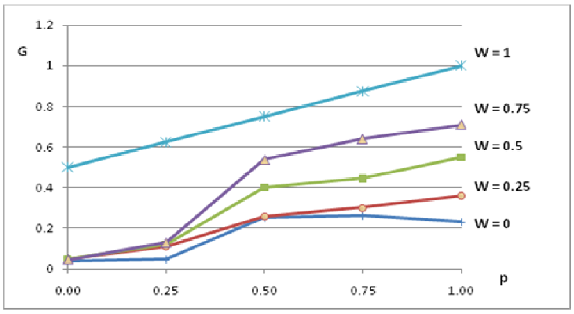

FIGURE 8. Dependence of the Equilibrium Value of G on p for Some Values of W

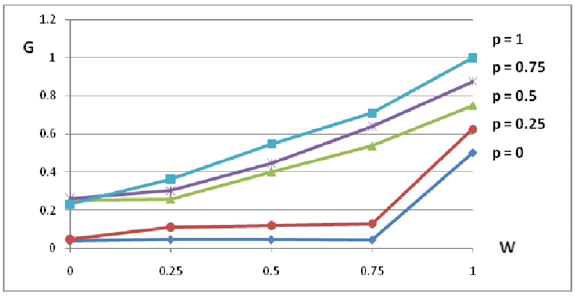

FIGURE 9. Dependence of the Equilibrium Value of G on W for Some Values of p

In Figure 8 it is shown that practically always the equilibria values for G depend on the basic personal characteristic p in a monotonically increasing manner; this expresses that the personal characteristic p indeed relates to the greed of the person, even while the latter is dynamic. Moreover, in Figure 9 it is shown that the equilibria values for G also depend on the state W of the economy in a monotonically increasing manner. This shows that a better economy leads to higher greed. Note, however, that for low values of p (from 0 to 0.3) the state W of the economy has not much impact, unless it is extremely high (W close to 1).

VI. Adapting the Risk Profile to the User

The risk profile *p* is a characteristic that depends on the person. As the model is to be applied to real humans, a mechanism is needed to estimate this value at runtime in an adaptive manner for a particular individual. To this end, this estimation is based on a person's decisions made. In each iteration, the model compares the actual product as selected by the human (X_h, Y_h) with the product that the software agent predicted to be selected (X_a, Y_a). For each of these products (X, Y) the value of $r(X, Y)$ is determined as indicated in Section 3. Then the value for *p* is adapted as follows.

$$p(t+\Delta t) = p(t) + \eta\, p(t)\, [(r(X_a, Y_a) - r(X_h, Y_h))/ r(X_h, Y_h)]\, \Delta t \quad \text{if} \quad r(X_a, Y_a) - r(X_h, Y_h) \leq 0$$

$$p(t+\Delta t) = p(t) + \eta\, (1-p(t))[(r(X_a, Y_a) - r(X_h, Y_h))/ r(X_h, Y_h)]\, \Delta t \quad \text{if} \quad r(X_a, Y_a) - r(X_h, Y_h) \geq 0$$

Here the adaptation rate η has been chosen *0.8*. Thus, this mechanism adapts the value for *p* with a percentage that is proportional to the difference between the actual and predicted value of *r*.

Some simulation experiments have been performed to test whether the agent correctly learns the risk profile of a (simulated) person. To this end, two versions of the decision making model were simulated: one for the human, and one for the agent which has to estimate the human's risk profile. Both models used the same formulae and parameter settings, but the values for *p* were different. For the human, different (static) values for *p* were taken. In all cases the model was able to find the person's risk profile in an accurate manner relatively quickly (at most within about 50 iterations). In some cases the model over-compensates a bit due to the high value of η, but this effect is minimal. Despite this positive result, one should be aware that these were ideal scenarios for the model to learn the human's risk profile, since the (simulated) humans behaved exactly according to the decision making model. For real humans, this behaviour will be more irregular, and thus the risk profiles will be more difficult to learn.

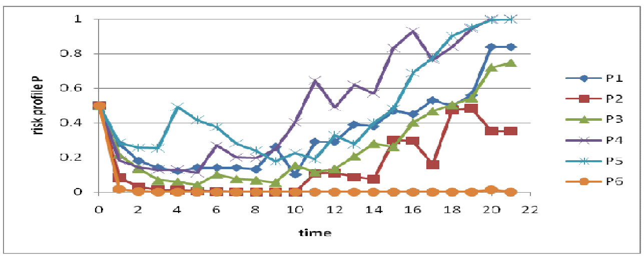

FIGURE 10. Risk Profile Adaptation for 6 Different Persons

During experiments conducted with humans that had to make a number of consecutive investment decisions (see [3] for more details), the values of the selected products, the person's estimated risk profile *p*, and greed value *G* have been logged.

This allows investigating the dynamics of the person's characteristics as estimated by the agent. The risk profile *p* of six different persons as learned by the agent during the experiment is shown in Figure 10. Here time is along the x-axis whereas the value of the estimated risk profile *p* is along the y-axis. As an example, person P5 is a person with an extremely high risk profile, as this person was always selecting products with a high risk value. As can be seen from Figure 10, initially the value of estimated value of *p* fluctuates a lot, but eventually this value approaches 1. Person P6 is a person with an extremely low risk profile. In this scenario, the person was always selecting products with low risk. As can be seen from Figure 10, initially the value of risk profile *p* drops to 0.017, then to almost 0 and stays around this value for the rest of the scenario. Moreover, as shown in Figure 10, P2 is moderately conservative, P1 and P3 are moderately aggressive, and both P4 and P5 are very aggressive investors.

In the above experiments, the users only knew the results of previous investments, but did not have knowledge about the state of the world economy while making their decisions. To test whether such knowledge would make a difference, some additional experiments (with 6 new participants) have been conducted, in which the users are made more aware of the economic situation during the scenario (i.e., at each round the value of *W*, in the domain [*0*, *1*], is shown to the user).

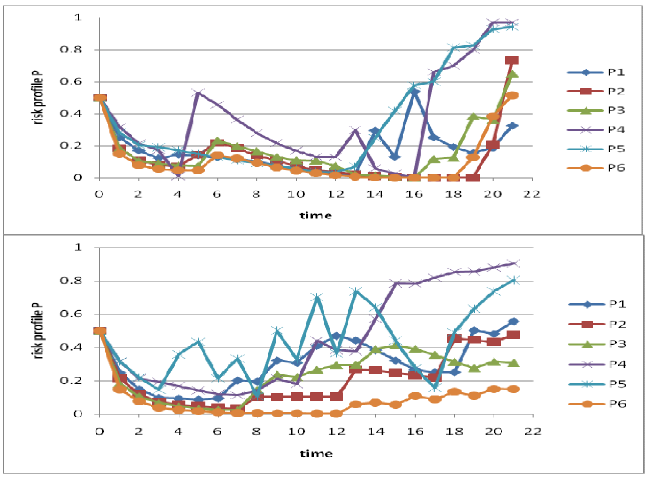

FIGURE 11. Risk Profile During two Scenarios (Fluctuating World Economy and Constant World Economy with $W = 0.5$)

Two different scenarios for the experiments have been created, one with a fluctuating value for *W*, and one in which *W* was constantly set to *0.5*. The results of these experiments are summarised in Figure 11. In the latter case, with a constant world state the values seem to stabilise more than in the other cases. These results may suggest that the person's level of awareness of the world state may be a nontrivial factor in this context.

VII. DISCUSSION

This paper presents an agent-based model of human decision making behaviour in economic situations, based on psychological states and characteristics concerning greed and risk taking or risk avoidance. In the area of modelling financial decision making within economics it is more and more acknowledged that psychological states and personality characteristics play an important role [4, 10, 12, 14]. Examples of such states are feeling insecure in relation to financial risks, and being greedy in relation to opportunities to obtain serious gains. The proposed agent model provides a computational formalisation of such concepts. The model takes ideas underlying the Modern Portfolio Theory [5, 13] into account, and incorporates the psychological concept greed and a risk characteristic. Thus a model is obtained that may provide a basis for the development of personalised intelligent agents that support a person in financial decisions.

To evaluate the model a number of simulation experiments have been performed, which illustrate the model's ability to mimic investment behaviour depending on the types of personality and the state of the economy. A mathematical analysis was contributed of the equilibrium states of the model. Moreover, some experiments have been conducted in which the model estimated the risk profiles of humans.

In recent years, various authors have studied processes in economics using agent-based simulation techniques (e.g., [2, 15, 16]. However, the current paper differs from these approaches in that it attempts to study the decision making behaviour of a single agent in detail, instead of analysing the global dynamics of a multi-agent system.

Future work aims at validation of the model, and at development of a virtual agent (avatar) [11] that supports the person by indicating choices of products that fit to the person's states and personality, but also by mirroring these aspects and implied behaviour in order to stimulate reflection by the person; a first step has been reported in [3]. As pointed out by various authors (e.g., [1]), virtual characters with personalisation features can augment involvement of users in financial (e.g., banking) applications.

REFERENCES

[1] Blom, J. and Monk, A. (2001). One-to-one e-commerce: who's the one? In: *Proceedings of the Conference on Human Factors in Computing Systems (CHI'01).* ACM Press, pp. 341-342.

[2] Bosse, T., Siddiqui, G.F., and Treur, J. (2010). Modelling Greed of Agents in Economical Context. In: N. García-Pedrajas et al. (eds.), *Proceedings of the Twenty Third International Conference on Industrial, Engineering & Other Applications of Applied Intelligent Systems, IEA-AIE'10, Part II*. Lecture Notes in Artificial Intelligence, vol. 6097. Springer Verlag, 2010, pp. 407-417.

[3] Bosse, T., Siddiqui, G.F., and Treur, J. (2010). An Intelligent Virtual Agents Increasing Involvement in Financial Services. In: *Proceedings of the 10th International Conference on Intelligent Virtual Agents, IVA'10*. Lecture Notes in Computer Science, Springer Verlag, 2010, to appear.

[4] Brandstätter, H. (1993). Should Economic Psychology Care about Personality Structure? *Journal of Economic Psychology*, vol.14, pp. 473-494.

[5] Haugen, R.A. (1997). *Modern Investment Theory*. Prentice-Hall, Inc.

[6] Kahneman, D. and Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, XLVII, pp. 263-291.

[7] Kasser, T., Cohn, S., Kanner, A., Ryan, R. (2007). Some costs of American corporate capitalism: A psychological exploration of value and goal conflicts. *Psychological Inquiry*, vol.18, pp.1-22.

[8] Kelley, C.T. (2003). Solving Nonlinear Equations with Newton's Method. *Fundamentals of Algorithms*, vol. 1. SIAM, Philadelphia, PA.

[9] Krawiec, K.D. (2000). Accounting for Greed: Unravelling the Rogue Trader Mystery. *Oregon Law Review*, vol 79, issue 2, pp. 301-339.

[10] Lo, A., Repin, D.V., and Steenbarger, B.N. (2005). Fear and Greed in Financial Markets: A Clinical Study of Day-Traders. *American Economic Review* 95, pp. 352-359.

[11] Prendinger, H., Lester, J., and Ishizuka, M. (eds.) (2008). *Intelligent Virtual Agents. Proceedings of the 8th International Conference on Intelligent Virtual Agents, IVA'08*. Springer LNAI, vol. 5208.

[12] Rabin, M. (2002). A Perspective on Psychology and Economics. *European Economic Review*, vol. 46, pp. 657-685.

[13] Sabal, J. (2002). *Financial Decisions in Emerging Markets*. Oxford University Press, Inc., New York.

[14] Simon, H.A. (1987). Behavioural Economics. In: *The New Palgrave: A Dictionary of Economics*. London, MacMillan.

[15] Terano, T., Deguchi, H., and Takadama, K. (eds.) (2003). Meeting the Challenge of Social Problems via Agent-Based Simulation. *Post-Proceedings of the 2nd International Workshop on Agent-Based Approaches in Economic and Social Complex Systems*. Springer Verlag, 2003.

[16] Tesfatsion, L. (2002). Agent-based computational economics: Growing economies from the bottom up. *Artificial Life*, vol. 8, pp. 55-82.

[17] http://www.ers.usda.gov/Data/Macroeconomics/