Micro to Macro Game Theory in a Multi-Agent System

Michael Harré
School of Information Technologies and
The Centre for the Mind
University of Sydney
Sydney, NSW, Australia
e-mail: mike@centreforthemind.com

Abstract

Recently there has been interest in the study of Multi-Agent Systems where these systems can be described using the principles of game-theory. In particular the Minority Game has been used to study the relationship between trading behaviour and volatility of financial markets while micro-economic games have been studied in order to understand how agents adapt to each others behaviour at the level of agent-agent interactions. This paper introduces a model which enables the numerical study of the collective adaptation of agents when a parameter ϕ is adjusted, varying the relative contribution of local and global interactions to the adaptive dynamics of the agents. The model also allows for the variation in a memory parameter which attenuates an agents history. It is demonstrated that for two different types of local interaction between the agents, co-operative and non-co-operative, there are significant memory dependent differences in the volatility of the system.

1 Introduction

In the work of Sato *et.al* [1] an agent is defined in terms of its action space, the history (memory) of rewards received from those actions and a selection procedure for mapping past {action, reward} pairs to an action that will be taken in the next time interval. To this end, and motivated by the literature on reinforcement learning [2, 3, 4], the principal idea is to define an agent $i \in \{1, 2, \ldots, N\}$ called σ_i which at time $\tau \in \{0, 1, 2, \ldots\}$ is required to choose an action j from a selection of M mutually exclusive alternatives, we label an agent so defined $\sigma_i^j(\tau)$. The agent has a state vector associated with it, called its choice distribution, of length M which reflects the probability associated with the agent taking action $j \in \{1, \ldots, M\}$ at time τ: $\mathbf{p}(\tau) = \{p(\sigma_i^1(\tau)), p(\sigma_i^2(\tau)), \ldots, p(\sigma_i^M(\tau))\}$ where $\sum_{j=1}^{M} p(\sigma_i^j(\tau)) = 1$. The agents we will study here will only have two possible action states to choose from, so we label each agent's action as either $\sigma_i^{\uparrow}(\tau)$ or $\sigma_i^{\downarrow}(\tau)$. Now we define $r_j(\tau)$ as the feedback an agent receives from the environment at time τ for making move $j \in \{\uparrow, \downarrow\}$ and we also define the update of the memory of past rewards the agent has received as [5]:

$$U_i^j(\tau+1) - U_i^j(\tau) = \frac{1}{T}\left[\delta_j(\tau) r_j(\tau) - \alpha U_i^j(\tau)\right] \quad (1)$$

where $U_i(0) = 0$, T is a constant that sets the agent-environment interaction time scale[1], $\alpha \in [0, 1]$ is a memory term ($\alpha = 0$ represents perfect memory of past feedback whereas $\alpha = 1$ represents no memory of past feedback and the agent's actions are instead entirely dependent on the most recent reward) and:

$$\delta_j(\tau) = \begin{cases} 1, & \text{action } j \text{ is chosen at time } \tau \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

Note that in the work on the Minority Game in [6] $\alpha = 0$, *i.e.* perfect memory of past actions, is implicit in the update of $U_i(\tau)$. Now we can define the probability distribution over the j possible responses at time τ as:

$$p(\sigma_i^j(\tau)) = \frac{e^{\beta U_i^j(\tau)}}{\sum_{k \in \{\uparrow,\downarrow\}} e^{\beta U_i^k(\tau)}}, \quad j \in \{\uparrow, \downarrow\} \quad (3)$$

In the work of Wolpert *et.al.* [7, 8] the β value is used to describe the agent's 'rationality', *i.e.* for $\beta \to \infty$ the agent will converge on the action with the greatest probability associated with it (totally rational) and for $\beta = 0$ the agent will act randomly without consideration of past rewards (totally irrational).

[1] we will assume $T = 1$ for the rest of this paper

International Conference on Computational Intelligence for Modelling
Control and Automation,and International Conference on
Intelligent Agents,Web Technologies and Internet Commerce (CIMCA-IAWTIC'06)
0-7695-2731-0/06 $20.00

IEEE COMPUTER SOCIETY

This work presents the initial findings of a particular structure of $U_i^j(\tau)$ based on the microscopic interactions of a version of the game of 'matching pennies' [1, 5, 9] and the macroscopic interactions of the Minority Game [10, 11, 12, 13]. In the microscopic interactions we model each agent σ_i as interacting with the two agents, σ_{i-1} and σ_{i+1}, on either side in a chain through a 2×2 interaction matrix B. The 'reward' for this local interaction with respect to σ_i is given by:

$$r^{local} = \sigma_{i-1} \cdot B \cdot \sigma_i + \sigma_{i+1} \cdot B \cdot \sigma_i \quad (4)$$

where $x \cdot y$ denotes the standard dot product between x and y and we assume that the following holds for any σ: $\sigma \in \{\sigma^\uparrow, \sigma^\downarrow\} = \left\{\left[\begin{smallmatrix}1\\0\end{smallmatrix}\right], \left[\begin{smallmatrix}0\\1\end{smallmatrix}\right]\right\}$. The matching pennies game uses:

$$B = \begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix}$$

so r^{local} is maximised for $\sigma_{i-1} = \sigma_{i+1} \neq \sigma_i$. In the Minority Game with zero information [13] we can describe the reward a single agent receives for its action choice through a similar equation:

$$r^{global} = \overline{\sigma} \cdot B \cdot \sigma_i \quad (5)$$

$$\overline{\sigma} = \sum_{i=1}^{N} \frac{\sigma_i}{N} \quad (6)$$

Note that for $\overline{\sigma} = \left[\begin{smallmatrix}0.5\\0.5\end{smallmatrix}\right]$, *i.e.* there are an equal number of agents of both orientations, $r^{global} = 0$ irrespective of whether or not $\sigma_i = \sigma_i^\uparrow$ or $\sigma_i = \sigma_i^\downarrow$ and r^{global} is positive only if $\overline{\sigma} \cdot \sigma_i < \frac{1}{2}$, *i.e.* σ_i is in the minority.

The total value of the reward associated with a given agent's orientation is defined as:

$$r = \phi r^{local} + (1 - \phi) r^{global} \quad (7)$$

where ϕ is a real valued parameter in the closed unit interval, $\phi \in [0, 1]$, which represents the significance each agent associates with the local reward relative to the significance associated with the global reward. r is used to construct an agent's memory of past rewards via equation 1 such that U^j, the agents utility, is the attenuated sum of past rewards for action j. For $\alpha = 1 : U_i^j(\tau + 1) = \delta_j(\tau) r(\tau)$. In this paper we are interested in the dynamic evolution of this system with particular interest in the way local and global interactions influence the overall system dynamics.

In section 2 we consider the case of agents with no memory and the dynamics of this system for local and global interactions between the agents. Section 3 addresses agents with memories and section 4 discusses how different types of local interactions give rise to differing system dynamics. Conclusions are discussed in the last section.

2 Memory-less Agents

In the first instance we consider the case where the agents have no memory of their past rewards and their behaviour is based exclusively on their current orientation, *i.e.* $\alpha = 1$ in equation (1). At each time-step τ there is a choice that each agent has to make: stay in the current state or change to the second state. If an agent has no memory of past rewards, the only information available is the current reward and the fact that there is an alternative state to be in, the agents are unaware of the potential value of the alternative state and it therefore assigns a 'guess' of $r_{alt}(\tau) = 0$ whereas the reward associated with the current configuration is $r_{current}(\tau) = \phi r^{local} + (1 - \phi) r^{global}$. In order to understand what the likely dynamics of this system are, we need to understand how equation (3) evolves via equation (7) and so initially we examine the two extreme cases where ϕ is either 0 or 1.

2.1 The $\phi = 1$ Case

In the $\phi = 1$ case only the local interactions have any influence on the dynamics and we can enumerate all the possible configurations of local interactions. We can do this by considering a chain of agents: $\{\ldots, \sigma_{i-1}(\tau), \sigma_i(\tau), \sigma_{i+1}(\tau), \ldots\}$, the probability for the next state which $\sigma_i(\tau)$ will be in is a conditional probability, *i.e.*:

$$p\Big(\sigma_i^\uparrow \to \sigma_i^\uparrow \Big| \sigma_{i-1}^\downarrow, \sigma_{i+1}^\uparrow\Big)$$
$$= C^{-1} \exp \beta(\sigma_{i-1}^\downarrow \cdot B \cdot \sigma_i^\uparrow + \sigma_{i+1}^\uparrow \cdot B \cdot \sigma_i^\uparrow)$$
and :
$$C = \exp \beta(\sigma_{i-1}^\downarrow \cdot B \cdot \sigma_i^\uparrow + \sigma_{i+1}^\uparrow \cdot B \cdot \sigma_i^\uparrow) + 1$$

where the $+1$ term in C is due to $r_{alt} = 0 \to e^{r_{alt}} = 1$ discussed previously. Note that if we consider in the example above the probability of moving from the current state to the alternative state, we would simply have: $p\big(\sigma_i^\uparrow \to \sigma_i^\downarrow \big| \sigma_{i-1}^\downarrow, \sigma_{i+1}^\uparrow\big) = C^{-1}$ for the this same reason.

The following table is a list of the probabilities for transitions given the local environment for memory-less agents (up to symmetrically equivalent configurations):

International Conference on Computational Intelligence for Modelling
Control and Automation,and International Conference on
Intelligent Agents,Web Technologies and Internet Commerce (CIMCA-IAWTIC'06)
0-7695-2731-0/06 $20.00 © 2006 IEEE

Transition Type	Conditional		Probability
1: $\sigma_i^{\uparrow} \to \sigma_i^{\uparrow}$	$\sigma_{i-1}^{\uparrow}$	$\sigma_{i+1}^{\uparrow}$	$\frac{1}{1+e^{2\beta}} < \frac{1}{2}$
2: $\sigma_i^{\uparrow} \to \sigma_i^{\uparrow}$	$\sigma_{i-1}^{\downarrow}$	$\sigma_{i+1}^{\uparrow}$	$\frac{1}{2}$
3: $\sigma_i^{\uparrow} \to \sigma_i^{\uparrow}$	$\sigma_{i-1}^{\downarrow}$	$\sigma_{i+1}^{\downarrow}$	$\frac{1}{1+e^{-2\beta}} > \frac{1}{2}$
4: $\sigma_i^{\downarrow} \to \sigma_i^{\uparrow}$	$\sigma_{i-1}^{\uparrow}$	$\sigma_{i+1}^{\uparrow}$	$\frac{1}{1+e^{2\beta}} < \frac{1}{2}$
5: $\sigma_i^{\downarrow} \to \sigma_i^{\uparrow}$	$\sigma_{i-1}^{\downarrow}$	$\sigma_{i+1}^{\uparrow}$	$\frac{1}{2}$
6: $\sigma_i^{\downarrow} \to \sigma_i^{\uparrow}$	$\sigma_{i-1}^{\downarrow}$	$\sigma_{i+1}^{\downarrow}$	$\frac{1}{1+e^{-2\beta}} > \frac{1}{2}$

Before discussing the numerical results it should be noted that there is one 'absolutely' stable and one 'periodically' stable configuration each agent can be in. In the absolutely stable case we can have a chain of agents: $\{\ldots, \sigma_{i-2}^{\uparrow}, \sigma_{i-1}^{\downarrow}, \sigma_i^{\uparrow}, \sigma_{i+1}^{\downarrow}, \sigma_{i+2}^{\uparrow}, \ldots\}$ which as $\beta \to \infty$ has a zero probability of changing and is therefore temporally frozen. This configuration is the Nash equilibrium for the local interactions as there is no benefit to any agent in this configuration in changing its orientation. On the other hand the chain: $\{\ldots, \sigma_{i-2}^{\downarrow}, \sigma_{i-1}^{\downarrow}, \sigma_i^{\downarrow}, \sigma_{i+1}^{\downarrow}, \sigma_{i+2}^{\downarrow}, \ldots\}$ as $\beta \to \infty$ has a probability of 1 of changing to the opposite configuration in the next time step and then changing back again at the time-step after that. Consequently this chain of agents has a temporal periodicity. This is clearly not a Nash equilibrium because all of the agents will be in the same state all of the time, so no agent ever has a different penny from its two neighbours. The other configurations have a fixed probability of $\frac{1}{2}$ irrespective of the value of β.

With $\phi = 1$ there can be no phase transition for variation in β [14, 15]. This is because phase transitions in models such as this only exist in the case of infinite range interactions between the agents and $\phi = 1$ insures there are only local interactions. However it is interesting to note how the system evolves for increasing β as there are clear signs that the two types of stable behaviours discussed above emerge smoothly from the random behaviour of low β (left plot Table 1). In this graphic, 250 agents are arrayed across the top of the plot with initial conditions of $p(\sigma_i^{\uparrow}(0)) = p(\sigma_i^{\downarrow}(0)) = \frac{1}{2}$ and $\beta = 0$. The two possible configurations that the agents can be in are plotted as a white point or a black point. Each successive update of the system is then plotted below the previous and every 20^{th} time step the β value increases by 0.2, β reaches a maximum of 5 and the left most agent is connected to the right most agent so that the agent chain is connected in a loop. It can be seen that the two types of stable solutions discussed above emerge smoothly from the random phase of low β values and the boundaries between the two regions of different solution types are composed of transition types 2 and 5 from the previous table, *i.e.* they have an equal probability of being in either orientation and this probability is independent of β.

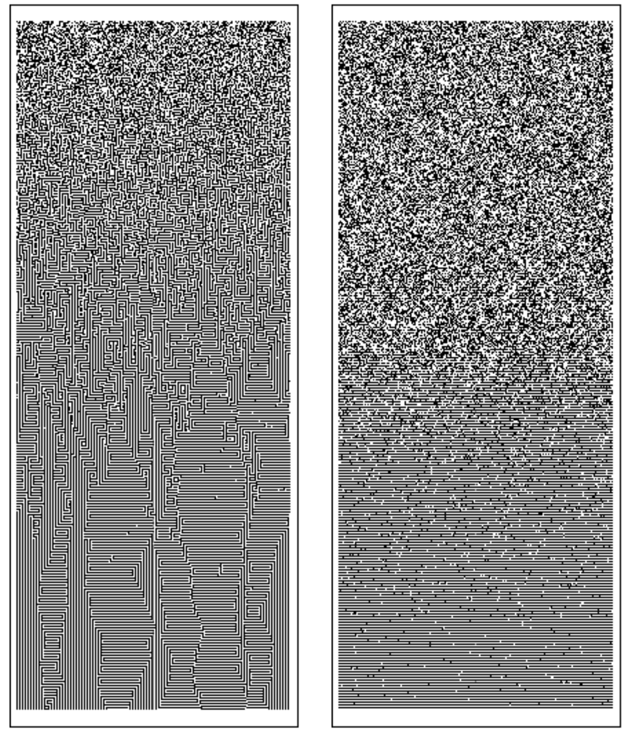

Table 1. Action orientations for 250 interacting agents (left: $\phi = 1$ (local), right: $\phi = 0$ (global)) with β increasing with τ from top to bottom. $\tau = \beta = 0 \to \tau = 500 = 100\beta$**.**

2.2 The $\phi = 0$ Case

In the $\phi = 0$ case (right plot, Table 1) we have the behaviour of the agents entirely dominated by the aggregate orientation of all agents in the system. A consequence of taking this average value is that $\overline{\sigma}$ in equation (5) is not binary but instead it can take on intermediate orientations which wouldn't exist except as an aggregate of other orientations. This leads to the natural instability of the configuration $\{\ldots, \sigma_{i-2}^{\uparrow}, \sigma_{i-1}^{\downarrow}, \sigma_i^{\uparrow}, \sigma_{i+1}^{\downarrow}, \sigma_{i+2}^{\uparrow}, \ldots\}$ discussed previously. In this case we see that the overall orientation of the system is:

$$\overline{\sigma} = \sum_{i=1}^{N} \frac{\sigma_i}{N} = \begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix} \tag{8}$$

and that $r^{global} = 0$ and so equation (7) is: $r = \phi r^{local} + (1 - \phi) r^{global} = 0$, *i.e.* no net reward

International Conference on Computational Intelligence for Modelling Control and Automation,and International Conference on Intelligent Agents,Web Technologies and Internet Commerce (CIMCA-IAWTIC'06)
0-7695-2731-0/06 $20.00 © 2006 IEEE

for this type of configuration and the agents will act randomly unless there is an alternative stable configuration for the agents to be in. There is an alternative available in the temporally periodic solution discussed above which is far from the Nash equilibrium, so this is the only solution we should expect, *i.e.* the periodic behaviour associated with the agent chain: $\{\ldots, \sigma_{i-2}^{\downarrow}, \sigma_{i-1}^{\downarrow}, \sigma_{i}^{\downarrow}, \sigma_{i+1}^{\downarrow}, \sigma_{i+2}^{\downarrow}, \ldots\}$ (right plot, Table 1). We should also expect to see a phase transition associated with variation in β as $\overline{\sigma}$ induces long range interactions between the agents (Figure 1). This phase transition in the Minority Game was analysed in [16] and marks the transition from low variability to high variability in the agent's behaviour at some critical value β_c. The oscillations from one configuration to another have a periodicity of 2 and it is this transition to cyclic behaviour which is responsible for the vertical striations we see dominating the lower half of the right hand plot in Table 1.

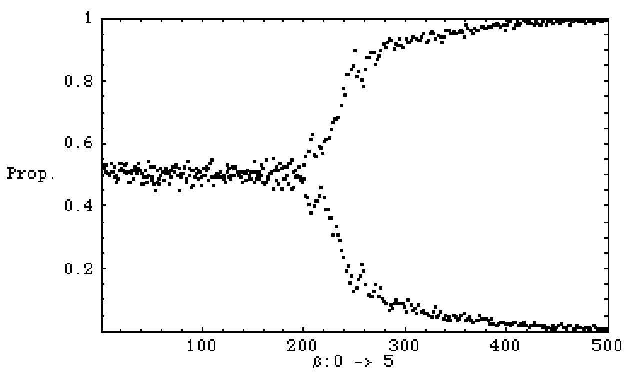

Figure 1. Proportion of total agents in one of the possible orientations for the right plot in Table 1 showing the phase transition at $\tau \approx 200$ ($\beta \approx 2$)

3 Agents With Memory

In [11, 17] extensive analytical and numerical studies were made of a Minority Game which is equivalent to $\phi = 0$ and $\alpha = 0$ *i.e.* perfect memory and no contribution to the dynamics from local interactions. In this section we present numerical results concerning the relationship between α and ϕ for some fixed $\beta > \beta_c$.

The examples we consider are for $\alpha = 0.6, 0.05$ and 0.0 where $\phi : 1 \rightarrow 0$. This allows us to see how the dynamics of the system evolve as it progresses from purely local to purely global dynamics. We start with the same initial conditions we had for the local dynamics in the previous example except we fix $\beta = 5$ and now for every 20^{th} update ϕ decreases by 0.1 until it reaches 0 and then the system is allowed to run on so that the dynamics for $\phi = 0$ can be observed. The rest of the parameters are set as they were for the previous example, 250 agents with random initial conditions arranged in a ring. The results are shown in Table 2 with ϕ decreasing from top to bottom, zero being reached half way down.

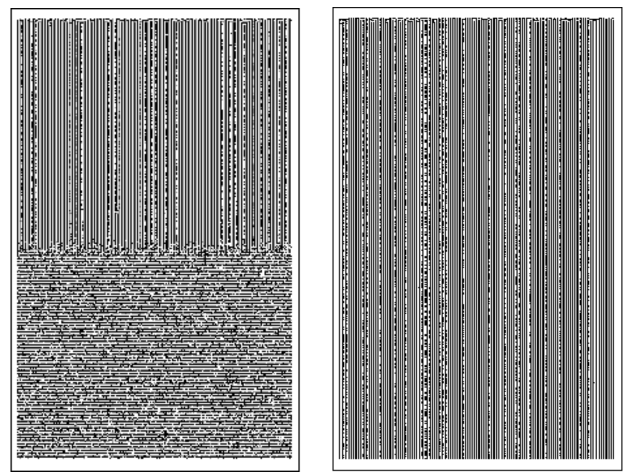

Table 2. Action orientations for $N = 250$ agents (left: $\alpha = 0.6$, right: $\alpha = 0.0$) with ϕ decreasing from top to bottom. $\phi = 0$ half way down.

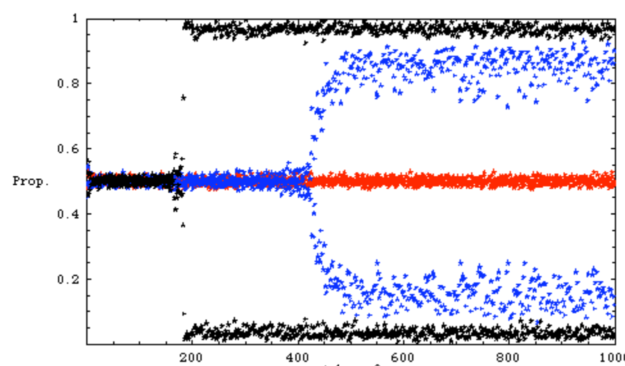

Figure 2. Relative proportion of total agent orientations for different values of α: black $\alpha = 0.6$ (branch at $\tau \approx 200, \phi = 0$), blue $\alpha = 0.05$ (branch at $\tau \approx 400$) and red $\alpha = 0.0$ (no branch)

It can be seen that before $\phi = 0$ the agents with fixed

International Conference on Computational Intelligence for Modelling Control and Automation,and International Conference on Intelligent Agents,Web Technologies and Internet Commerce (CIMCA-IAWTIC'06)
0-7695-2731-0/06 $20.00 © 2006 IEEE

positions dominate with a small proportion regularly switching their orientation. Just prior to $\phi = 0$ there is a brief random interval before the dynamics flip and the system is dominated by the period two oscillatory behaviour associated with $\beta > \beta_c$. As we decrease α the variance in the region $\phi = 0$ decreases very slowly with α but not until $\alpha = 0$ do we see the variance in the $\phi = 0$ region achieve the same low value as is seen in the $\phi < 0$ region. Note that there is very little volatility and very little dependence on α in the region $\phi > 0$ compared to $\phi = 0$. The point at which the oscillations (branch points in Figure 2) set in also increases with increasing memory. As numerical simulations show that this does not occur if we start the system with $\phi = 0$ and fixed, it is clear that this time delay to the oscillatory solution is due to a 'priming' effect induced by the agent's memory of good choices (high utility) in the region $\phi > 0$ ($\tau < 200$ in Figure 2). This is supported by the right hand plot in Table 2 where, for $\alpha = 0$, those agents which were fixed prior to $\phi = 0$ remain fixed after $\phi = 0$.

4 Agent Co-operation Reduces Volatility

As mentioned above, the dynamics for this system in the $\alpha = 0$, $\phi = 0$ state have been well studied in the past and the issue of reducing the volatility in the agent's behaviour has often been addressed for example by introducing more information into the model such as 'public' information [10, 11]. However, it has also been suggested that volatility can be reduced by adding co-ordination to the dynamic interactions of the agents [13]. Low volatility is desirable in the global dynamics of the agents because it has been recognised that certain types of volatility, such as all agents switching from one configuration to the alternative configuration (lower half of left hand figure, Table 2) is far from the Nash equilibrium. In this case there would be no agents in the minority and therefore no agents earn any reward. We now present a way in which co-operation can be introduced into the system which improves co-ordination between the agents, such as coalition block formation, which enables volatility to be reduced significantly.

We recall the structure of equation 4 where the local interactions were based on the local equivalent of the Minority Game's global dynamics, called the matching pennies game. If instead of B being used for both the local and global dynamics we replaced B in r^{local} with another matrix A:

$$A = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$$

so that now $r^{local} = \sigma_{i-1} \cdot A \cdot \sigma_i + \sigma_{i+1} \cdot A \cdot \sigma_i$. This payoff matrix implies a co-operative relationship between the agents such that r^{local} is maximised when $\sigma_{i-1} = \sigma_i = \sigma_{i+1}$. In order to measure the volatility reduction 250 agents were modelled interacting over 125 time-steps with $\beta = 5$ and $\phi = 0.1$ so that there was only a 10% contribution to the overall dynamics coming from the local interactions, the other 90% coming from the global Minority Game interactions. In Figure 3 we show the results in the volatility reduction (measured as variance) as a function of decreasing memory for both non-co-operative and co-operative local dynamics. Most of the volatility has been reduced for locally co-operative agents for $\alpha < 0.5$, a sharp contrast with the the dynamics of locally non-co-operative agents which show high volatility for $\alpha > 0.25$ where below this value both methods are very similar in their results. It is also interesting to see that even with zero memory ($\alpha = 1$), local co-operation helps reduce variability.

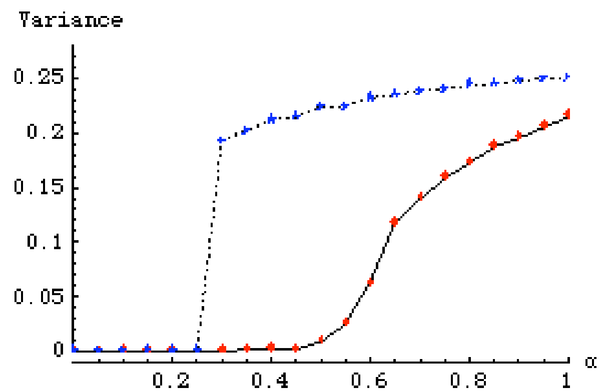

Figure 3. Variance in the agent behaviour as a function of memory ($\phi = 0.1$, $\beta = 5$) for non-co-operative (blue, top curve) and co-operative agents (red, bottom curve).

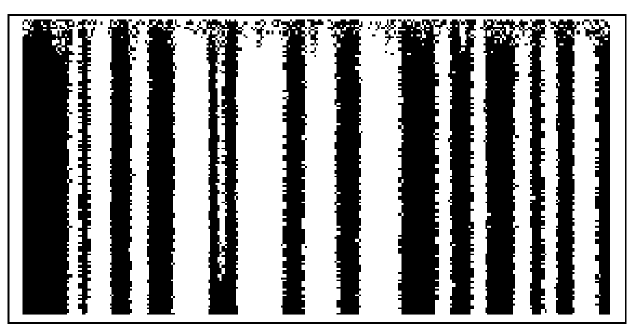

Figure 4. Agent orientations for $N = 250$, $\beta = 5$, $\phi = 0.1$, $\alpha = 0.1$ and $\tau : 1 \rightarrow 125$.

International Conference on Computational Intelligence for Modelling Control and Automation,and International Conference on Intelligent Agents,Web Technologies and Internet Commerce (CIMCA-IAWTIC'06)
0-7695-2731-0/06 $20.00 © 2006 IEEE

We can see what is happening by again looking at a plot of the agent orientations for locally co-operative agents as a function of τ (Figure 4). After an initial transient phase the agents form aggregates of agents of the same orientation which are mostly fixed in their behaviour. There are however agents which form the boundaries between these coalition blocks which often exhibit switching behaviour and these are the agents which contribute most of the variability to the system. It has been pointed out [13] that frozen agents are more efficient than oscillating agents in the Minority Game and by adding a degree of local co-ordination between the agents it is this frozen behaviour we recognise as the source of reduced volatility in the system.

5 Conclusion

In this paper we have introduced a model which enables us to study the relationships between local and global dynamics for agents interacting in a game theoretical setting. This approach is particularly interesting as many real world scenarios are not dominated solely by the local or the global game dynamics, but have components of both contributing to the performance of the agents and the system as a whole. Several numerical results have been presented demonstrating the rich range of dynamics this system exhibits when there are interactions on both a micro and a macro level. Most importantly it was shown that there is a non-trivial dependency of the system dynamics on the attenuation of the agents memory which, depending on the nature of the local interactions, can have an impact on the overall volatility in the agents behaviour.

6 Acknowledgements

The author would like to thank Terry Bossomaier and Allan Snyder for their invaluable discussions and continuing support. This work was supported in part by ARC grant number DP0560207.

References

[1] Y. Sato, E. Akiyama, and J. Farmer. Chaos in learning a simple two-person game. *PNAS*, Vol. 99 number 7:4748 – 4751, 2002.

[2] T. Börgers and R. Sarin. Learning through reinforcement and replicator dynamics. *Journal of Economic Theory*, Vol. 77, issue 1:1 – 14, 1997.

[3] L. Kaelbling, M. Littman, and A. Moore. Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, Vol. 4:237 – 285, 1995.

[4] R. S. Sutton and G. A. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 1998.

[5] Y. Sato, E. Akiyama, and J. Crutchfield. Stability and diversity in collective adaptation. *arXiv: cond-mat/9907176*, 1999.

[6] Marsili M. Zecchina R. Challet, M. Statistical mechanics of systems with heterogeneous agents: Minority games. *Phys. Rev. Lett.*, Vol. 84(8):1824 – 1827, 2000.

[7] D. Wolpert. Information theory: The bridge connecting bounded rational game theory and statistical physics. In D. Braha and Y. Bar-Yam, editors, *Complex Engineering Systems*. Perseus Books, 2004.

[8] D. Wolpert. A predictive theory of games. *arXiv: nlin.AO/0512015 v1*.

[9] Y. Sato and J. Crutchfield. Coupled replicator equations for the dynamics of learning in multiagent systems. *arXiv:nlin.AO/0204057*, 2002.

[10] Zhang Y. Challet, D. Emergence of cooperation and organization in an evolutionary game. *Physica A*, Vol. 246:407 – 408, 1997.

[11] Zhang Y. Challet, D. On the minority game : Analytical and numerical studies. *Physica A*, Vol. 256:514 – 532, 1998.

[12] Challet D. Zecchina R. Marsili, M. Exact solution of a modified el farol's bar problem: Efficiency and the role of market impact. *Physica A*, 3 - 4:522 – 553, 2000.

[13] Challet D. Marsili, M. Trading behavior and excess volatility in toy markets. *Advances in Complex Systems*, Vol. 4(1):3 – 17, 2006.

[14] F. Dyson. Esistence of a phase-transition in a one-dimensional ising ferromagnet. *Commun. math. Phys.*, Vol. 12:91 – 107, 1969.

[15] F. Dyson. An ising ferromagnet with discontinuous long-range order. *Commun. math. Phys.*, Vol. 21:269 – 283, 1971.

[16] R. Savit, R. Manuca, and R. Riolo. Adaptive competition, market efficiency, and phase transitions. *Phys. Rev. Lett.*, Vol. 82:2203 – 2206, 1999.

[17] D. Challet and Y. Zhang. Emergence of cooperation and organization in an evolutionary game. *Physica A*, 246 Number 3:407–418, 1997.

International Conference on Computational Intelligence for Modelling Control and Automation,and International Conference on Intelligent Agents,Web Technologies and Internet Commerce (CIMCA-IAWTIC'06)
0-7695-2731-0/06 $20.00