

Searching for good strategies in adaptive minority games

Marko Sysi-Aho, Anirban Chakraborti,* and Kimmo Kaski

Laboratory of Computational Engineering, Helsinki University of Technology, P.O. Box 9203, FIN-02015 HUT, Finland

(Received 13 May 2003; revised manuscript received 12 September 2003; published 31 March 2004)

In this paper we introduce and study various adaptive minority game models in which agents try to improve their performances by modifying their strategies through genetic algorithm based crossover mechanism. One aim of this study is to find out what happens at the system as well as at the individual agent level. Adaptation is found to improve the performance of individual agents quite remarkably, to tighten the competition among the agents, and to drive the whole system towards maximum efficiency. Results from four adaptive minority games and the basic minority game are compared, and the parameter dependencies of the best performing game are discussed.

DOI: 10.1103/PhysRevE.69.036125

PACS number(s): 02.50.Le, 87.23.Ge, 87.23.Kg

I. INTRODUCTION

Various systems of natural and societal origin show a complex behavior, which can be attributed to a competition among interacting agents for scarce resources and their adaptation to continuously changing environment [1–5]. Such agents could be diverse in form, function, and capability, for example, cells in an immune system or firms in a financial market. The nature of interactions between agents is dependent on their capabilities, and the behavior of an agent can be considered as a collection of rules governing *responses to stimuli*. In order to model these systems, a major concern is the selection and representation of the stimuli and responses, through which the behavior and strategies of the agents are defined. In a model, the rules of action serve as a direct way to describe the strategies of agents, and their behavior is studied by monitoring the effect of rules acting sequentially. The other key process to be included in the model is *adaptation*, which in biology serves as a mechanism for an organism to try to make itself fit to changing environment. What makes these systems fascinatingly complex is the fact that the environment of a particular agent includes other adaptive agents, all competing with each others. Thus, a considerable amount of an agent's effort goes into adaptation and reaction to the other agents. This is the main source of interesting temporal patterns and emergent behavior these systems produce.

In this paper, we will study a simple game model, in which agents adapt dynamically to compete and perform better. In such a model the strategies, which an agent uses to decide the course of action, must be very good or best for the agents to survive—similar to the “survival of the fittest” principle in biology. So just as an organism adapts itself to its natural environment, we propose that the agents of a game adapt themselves by modifying their strategies from time to time, depending on their current performances. For this purpose we borrow the concept of genetic crossover from biology and use it to modify the strategies of agents in the course of the game, in the same way as in genetic algorithms [6–8]. More specifically we apply this adaptation scheme to the

minority game, introduced in Refs. [9–13]. Although the behavior of the minority game is believed to expose a number of important characteristics of complex evolving systems, one of its weaknesses is that agents have limited possibilities to improve their own performance whereas in real competitive environment attempts to improve ones skills continuously are imperative. Our adaptation scheme [14,15] proposes a natural and simple way to take this essential feature into account, and its application turns out to give results quite different from those of the basic minority game and its variants [9,16–18].

This paper is organized such that next we briefly introduce our minority game model together with various genetic algorithm based adaptation mechanisms for strategy changes. This is followed by Sec. III, where we first compare comprehensive computer simulation results of these adaptive games and then analyze the parameter dependencies of the best performing game. Finally we draw conclusions.

II. MODEL

Let us first describe the basic minority game (BG) model of Challet and co-workers [9,10]. There one assumes an odd number of agents N which can perform one of two possible actions denoted here by 0 or 1. An agent wins a round of a game if it chooses the action belonging to the minority group. All the agents are assumed to have access to a finite amount of “global” information, in the form of a common bit-string “memory” of M most recent outcomes of the game, such that there are 2^M possible “history” bit strings. An agent's “strategy” consists of two possible responses, i.e., an action 0 or an action 1, to each possible history bit string. Thus, there are 2^{2^M} possible strategies constituting the whole “strategy space” Ω , from which each agent picks S strategies at random to form its own pool Ω_i , where $i = 1, \dots, N$ denotes an agent number. Each time the game is played, time t is incremented by unity and one “virtual” point is assigned to the strategies that have predicted the correct outcome and the best strategy is the one which has the highest virtual point score. An agent's performance is measured by the number of times the agent wins, and the strategy, which the agent used to win, gets a “real” point. The number of agents who choose one particular action changes with time, and is denoted by x_t .

In order to describe the collective behavior of the agents,

*Present address: Department of Physics, Brookhaven National Laboratory, Upton, NY 11973, USA.

we define the concept of scaled utility in terms of x_t , in the following way:

$$U(x_t) = \{[1 - \theta(x_t - x_M)]x_t + \theta(x_t - x_M)(N - x_t)\} / x_M, \quad (1)$$

where $x_M = (N - 1) / 2$ is the maximum number of agents who can win, and

$$\theta(x_t - x_M) = \begin{cases} 0 & \text{when } x_t \leq x_M \\ 1 & \text{when } x_t > x_M \end{cases}$$

is the Heaviside unit step function. When $x_t = x_M$ or $x_t = x_M + 1$, the scaled utility of the system is maximal, $U_{max} = 1$, as the highest number of agents win. The system is more efficient when deviations from the maximum utility U_{max} are small, or in other words, the fluctuations in x_t around the mean ($N/2$) are small.

At the level of individual agents, their performances in the basic minority game evolve such that the agents who begin to perform badly do not improve with time, and those who do well continue doing so [9,14]. This indicates that by chance well performing agents were blessed with good strategies while badly performing agents got bad strategies. However, there are competitive environments in which individual agents need to adapt themselves to do better or to survive. But being good at one moment does not guarantee that one would stay good later. In fact, there are many examples in business, sports, etc. which show that those who have decided to rest on their laurels have been superseded by those who have decided to adapt and fight back, persistently. This feature of dynamic competition needs to be included in the model, and it can be simply realized by allowing agents to modify strategies in their individual pools. How well an agent does then, in reality, depends on the agent's capabilities and skills, and how an agent refines its strategies.

For the adaptation or strategy modification we have chosen genetic algorithms [8], which have turned out to be useful in various optimization problems. Within the framework of the minority game the adaptation is realized by letting agents to check their performances after a time interval τ , and if an agent finds that it is among the worst performing fraction n (where $0 < n < 1$), it modifies its strategies by applying genetic operands to its strategy pool [14,15]. Here the quantity τ describes a time scale that characterizes the adaptation rate of agents in the system. Hence it can vary on a wide range for systems of natural origin to systems of societal nature.

In the genetic adaptation schemes of this study, an agent chooses two "parents" from its current pool of strategies $\Omega_i(t) \subset \Omega$, and draws a (uniformly distributed) random number to determine the crossover point. Then the parts of the strategies, above and below this point, are interchanged to produce two new strategies called the "offsprings." In addition there are various choices as for which strategies are selected as the parents and also which strategies are replaced by the reproduced offsprings. The mechanism which works the best depends on the circumstances and can vary from system to system. In some cases it is possible that saving the

parent strategies would threaten the success of the newborn strategies or create too stiff a competition among the strategies leading to possible disorder, and in other cases the opposite might happen. In this study, we have considered four different adaptation schemes by first selecting from the strategy pool of an agent, the parent strategies to perform a genetic one point crossover for reproducing offsprings, and then selecting two old strategies from the same pool to be substituted by the offspring strategies.

(A) Two parent strategies are selected at *random*, and after one-point genetic crossover the parent strategies are substituted with the two new strategies (offsprings).

(B) Two parent strategies are selected at *random* and after one-point genetic crossover the two worst performing strategies are substituted with the two new strategies (offsprings) while the parent strategies are saved.

(C) Two *best* performing strategies are selected as parents, and after one-point genetic crossover the parents are substituted with the two new strategies (offsprings).

(D) Two *best* performing strategies are selected as parents and after one-point genetic crossover two worst performing strategies are substituted with the two new strategies (offsprings) while the parent strategies are saved.

We would like to propose that these adaptation schemes could be considered in a loose sense to bear some resemblance to reality. From the point of view of choosing parents at random, schemes (A) and (B) correspond to "democratic" or equal opportunity reproduction, while schemes (C) and (D) are "elitist" due to selecting the best parents for reproduction. As for substitutions in the agents' strategy pools, in schemes (A) and (C) parents give space for their offsprings to live and develop without the need to fight with them for limited resources, a sacrifice for improving the survival of the species. Examples of parents dying after reproduction are numerous in nature. In decision making the interpretation of killing the parent strategies is that old strategies—unable to lead into success—are removed to give way to hopefully better strategies. Schemes (B) and (D), with parents being saved and agents getting rid of their worst strategies, bear some resemblance to "natural selection" of the fittest surviving species. On the other hand, in decision making situations these two schemes correspond to agents eradicating their losing strategies. Thus it is expected that schemes (B) and (D) lead to a tightening competition between agents. Furthermore, it could be expected that scheme (D) is the most efficient one, because it removes the worst strategies and replaces them with crossovers of the best ones, while saving the so far two best strategies in the game. In order to study the effects of the stiff competition between agents with continuously improving strategies in more detail, large scale simulations are needed. In these simulations it turns out that when agents use genetic operands, the scaled utility of the system increases and tends to maximize with different rates depending on the mechanism and the parameters of the game.

It should be noted that our genetic algorithm based mechanisms of evolution are considerably different from the mechanisms applied before within the framework of the minority games [9,16–18]. Here, the strategies are changed by

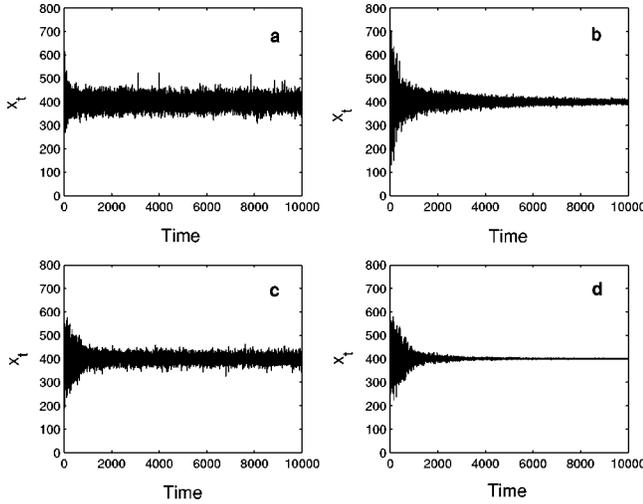


FIG. 1. Plots of x_t (the number of agents making a particular action) as a function of time, for the four adaptation mechanisms (A)–(D) described in Sec. II.

the agents themselves and they belong to the same strategy space Ω , that is not changing in size and dimension.

III. RESULTS

A. Comparison between adaptation mechanisms

In order to compare the above introduced four genetic adaptation mechanisms, we first study the quantity x_t , which describes the number of agents taking a particular action at a given time. For this we chose the parameters of the games as follows: $N=801$ agents with $M=6$ memories, $S=16$ strategies, and an adaptation time $\tau=40$ for the worst performing fraction $n=0.4$ of the agents, which we expect to be relevant in terms of the system size and obtainable statistics. Also we found out that in the game with adaptation scheme (D) the fluctuations increase monotonically with the control parameter $z=2^M/N$, i.e., without showing a phase change, unlike in the basic minority game. Thus for an even comparison the parameters were adjusted for the same phase in all the games. The simulation results are depicted in Fig. 1.

First, in Fig. 1(a) we present the results of adaptation scheme (A), in which the randomly chosen parent strategies are replaced with the reproduced offspring strategies. In this case it turned out that fluctuations in x_t around its mean (≈ 400) decay very rapidly from the initial level, which corresponds to the amount of fluctuations of the basic minority game, to a more or less constant level of less than half the initial value. This renders the scheme (A) game more efficient than the basic minority game. Second, in Fig. 1(b) we present the results of adaptation scheme (B), in which the offsprings strategies of randomly chosen parent strategies replace the two worst strategies in the agent's pool. In this case we observe that fluctuations in x_t around the mean decay, first rapidly below the value produced by scheme (A) and then slower to even smaller values. Thus the efficiency of the system is further improved. Third, in Fig. 1(c) we present the results of adaptation scheme (C), in which the best two strategies as parents are replaced after reproduction with their

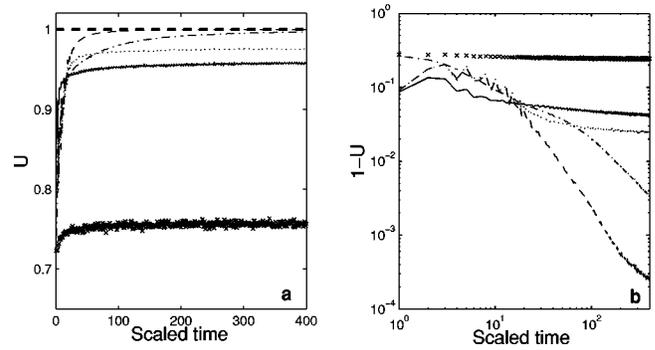


FIG. 2. (a) Scaled utility U and maximum utility $U_{max}=1$ (dashed horizontal line) and (b) utility deviation $(1-U)$ for four adaptation mechanisms as functions of scaled time (one unit of scaled time corresponds to a time average over a bin of 50 simulation time steps). Each curve is an ensemble average over 100 runs. In both panels, crosses represent the basic minority game, solid line represents adaptation mechanism (A), the dash dotted line represents mechanism (B), the dotted line represents mechanism (C), and the dashed line represents mechanism (D).

offspring strategies. In this case fluctuations in x_t around the mean once again decay rapidly then stabilizing to a level which is smaller than for adaptation scheme (A) but larger than for adaptation scheme (B). Fourth, in Fig. 1(d) we present the results of adaptation scheme (D), in which the offspring strategies of the best two strategies as parents replace the two worst strategies in the agent's pool. In this case we see that fluctuations in x_t die off very rapidly, thus making the system most efficient.

Next we investigate the scaled utility $U(x_t)$, defined in Eq. (1), which apart from the efficiencies of the games is expected to give insight to their dynamical behavior. Instead of the standard practice of studying the variation of σ^2/N versus $2^M/N$, where $|\sigma|$ stands for the difference in the number of agents between the majority and minority groups, we have studied U , because fluctuations in x_t decay strongly for adaptation mechanisms (B) and (D), in the latter case sometimes even disappearing completely. In Fig. 2 we show the results of (a) the scaled utility U and (b) the utility deviation $1-U$ from the maximum $U_{max}=1$ as a function of the scaled time for the four adaptation mechanisms, using the same set of parameters as before (i.e., $N=801$, $M=6$, $S=16$, $\tau=40$, and $n=0.4$). We find that the scaled utility rapidly saturates for the basic minority game, to efficiency level considerably less than the maximum. On the other hand, it is clearly seen that our four adaptation mechanisms greatly enhance the utilities close to the maximum. However, as evident in Fig. 2(b), the utility enhancement for mechanisms (A) and (C) seem to slow down or possibly even saturate to values slightly below $U_{max}=1$ while for mechanisms (B) and (D) U_{max} is approached more rapidly or possibly asymptotically. Of the latter two mechanisms the efficiency of scheme (D) game improves the fastest and is best overall. The same tendencies were also visible in Fig. 1. Thus we can conclude that the “elitist” adaptation scheme, in which only the best performing strategies can be used to reproduce new offspring strategies to replace worst performing strategies, is

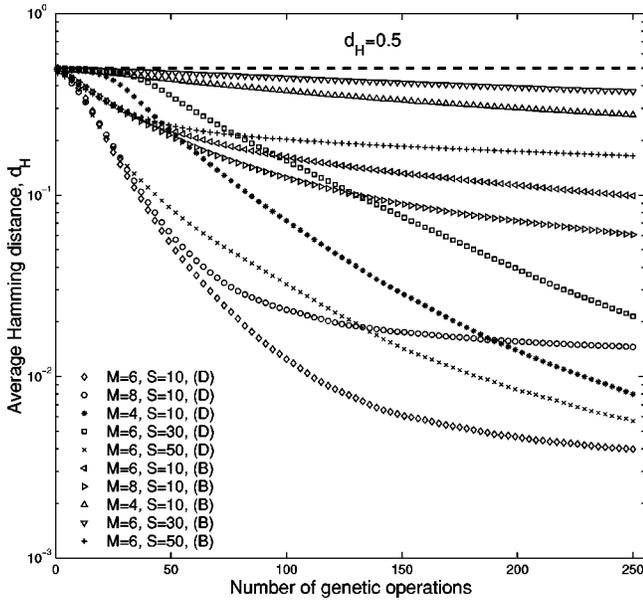


FIG. 3. The average Hamming distance vs the number of genetic operations (one genetic operation takes place after every τ time steps) for different adaptation mechanisms and parameters. The simulation was done with $N=801$, $\tau=80$, $n=0.4$ and for different combinations of memories M , strategies S , and adaptation schemes [(B) or (D)], as indicated in the legend. Each point is an ensemble average over 20 runs. The average Hamming distance $d_H=0.5$ is shown as a dashed horizontal line.

the best adaptive game. Later we will return to scheme (D) game to study in detail its parametric dependencies.

In order to examine the evolution of strategies in the agents' pools we use the Hamming distance, denoted by d_H . It measures how similar the strategies are and is defined between two strategies as the ratio of the uncommon bits to the total length of the strategy. The strategies are "correlated," if all the bits are pairwise the same, i.e., $d_H=0$, "anticorrelated" if all bits are opposite, i.e., $d_H=1$, and "uncorrelated" when exactly one half of the bits differ, i.e., $d_H=0.5$.

Here we will consider the average Hamming distance, which is calculated by first taking the average of the Hamming distances over all possible strategy pairs in the agent's pool, and then taking the average over all the agents. While it is obvious that individual Hamming distances between strategy pairs can change as a result of genetic crossovers, the situation is more complex for the average Hamming distance. As a matter of fact in the adaptation schemes in which the parent strategies after crossover are replaced by their offsprings [i.e., schemes (A) and (C)], all the bits in the agent's strategy pool and for that matter in the whole strategy space remain the same, and thus the average Hamming distance does not change. Therefore, this measure is useful only for the games where the bits in a strategy pool can change over time, i.e., adaptation schemes (B) and (D), for which the results are depicted in Fig. 3. Here it is seen that as the game evolves, the average Hamming distance decreases in both cases towards small values, but for scheme (B) game this happens considerably slower than for the scheme (D) game.

In the latter case d_H reaches a very small value, indicating that each agent tends to end up using a particular strategy in its pool for the best performance. In the case of scheme (B) the same seems to happen but it takes at least an order of magnitude longer time. On the other hand the plots of x_t (i.e., the number of agents choosing a particular action, depicted in Fig. 1) shows that these strategies are such that the total utility, and thus the efficiency of the system tends to maximize.

In Fig. 3 the results are shown for varying the memory size M and the number of strategies S in each agent's pool. We have observed that for adaptation mechanism (B) increasing M while keeping $S=10$ fixed makes the decay in the average Hamming distance faster, yielding $M=8$, the fastest decaying case. On the other hand, increasing S and keeping $M=6$ fixed does not seem to yield any systematic behavior, while the $S=10$ case seems to give rise to the fastest decay in the average Hamming distance. For adaptation mechanism (D) the situation is even less systematic, since increasing M and keeping $S=10$ fixed yields the fastest decaying $M=6$ case, and increasing S and keeping $M=6$ yields the fastest decaying $S=50$ case.

As a final comparison between different adaptation mechanisms we have studied a mixed game, in which agents can change their strategies with different adaptation schemes. In this game all the agents start by playing the basic minority game up to a given time $t=3120$ simulation time steps. Thereafter, three players continue playing with adaptation scheme (A) and another three with adaptation scheme (D), while the remaining players continue playing the basic minority game using their initially introduced strategy pools without adaptation. From the simulation results, depicted in Fig. 4, we find that the adaptive agents, some of which were the worst performing agents at the beginning, become quickly successful and outperform all the agents playing the basic minority game. As a matter of fact the slopes of the performance curves, i.e., the success rates of adaptive agents, is by far better than the best agent playing the basic minority game. In addition, we observe that all the agents using adaptation scheme (D) perform better than those using scheme (A), and that the competition between these three is tough. This serves as further evidence that adaptation scheme (D) is the most efficient one of the games discussed here, and thus interesting for a more detailed analysis.

B. Parametric studies

In this section we study the dependence of our adaptive game (D) on the model parameters, i.e., the memory size M , the number of strategies S in each agent's pool, the adaptation time τ , and the fraction of worst performing agents n . This is done in terms of the scaled utility U , defined in Eq. (1) and describing the efficiency of the game. Here we will use the quantity $1-U$ to illustrate the deviation from the maximum utility $U_{max}=1$. Below we present results of extensive simulations for utility deviation vs scaled time (one unit of scaled time corresponds to a time average over a bin of 50 simulation time steps) when the parameters were varied in pairs: (i) τ and M (Fig. 5), (ii) S and M (Fig. 6), and (iii) n and M (Fig. 7).

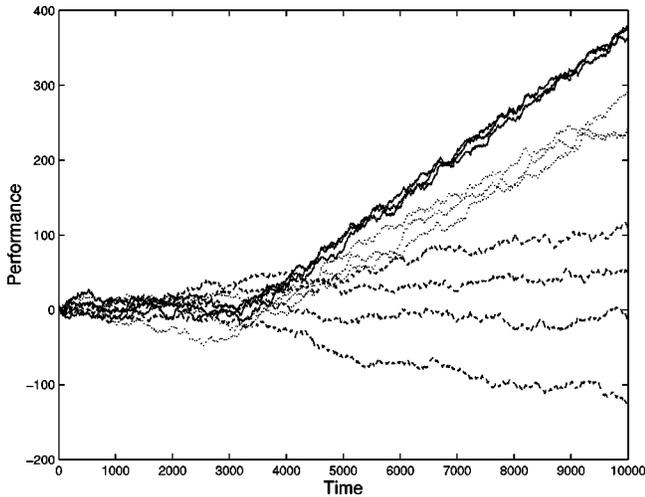


FIG. 4. Performances of selected agents as a function of time for a mixed game with agents having different adaptation schemes. The performances are scaled such that the mean performance of all the agents is zero. At $t=3120$, six agents begin to modify their strategies such that three of them (the three uppermost) used adaptation scheme (D) and the other three (the next three from top) used adaptation scheme (A). All the rest of the agents played the basic minority game without adapting; the performances of only the best, the worst, and two randomly chosen agents are shown (the four lowermost curves). Simulations were done with $N=801$, $M=8$, $S=16$, $n=0.3$, and $\tau=80$.

In Fig. 5 we show the effect of changing the adaptation time $\tau=10, 20, 40, 80$, and 160 (each panel) and memory $M=5, 6, 7$, and 8 (separate panels). For memory $M=5$ and 6 (two upper panels) and for all the studied adaptation times the utility deviation curves after initial transient periods are found to decrease; thus the games become more efficient. This transient appears to be directly proportional to the adaptation time, and the curves show more or less a linear decay region in scaled time, such that each region seems to

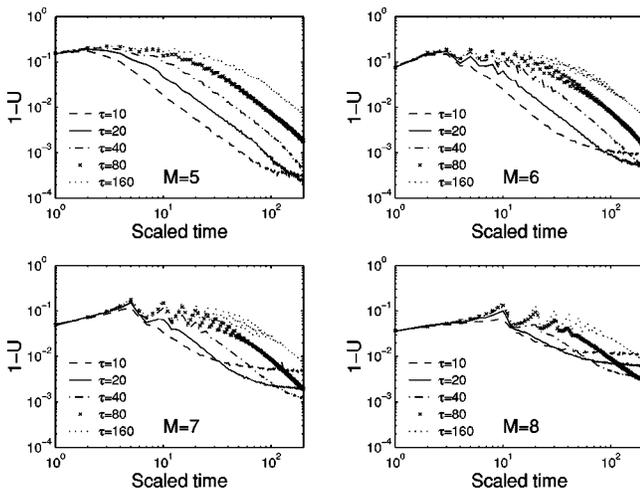


FIG. 5. Utility deviation $1-U$ vs scaled time: Effects of the adaptation time τ and memory M . Simulations were done with $N=1001$, $S=10$, and $n=0.3$, and ensemble averaging each curve over 50 runs.

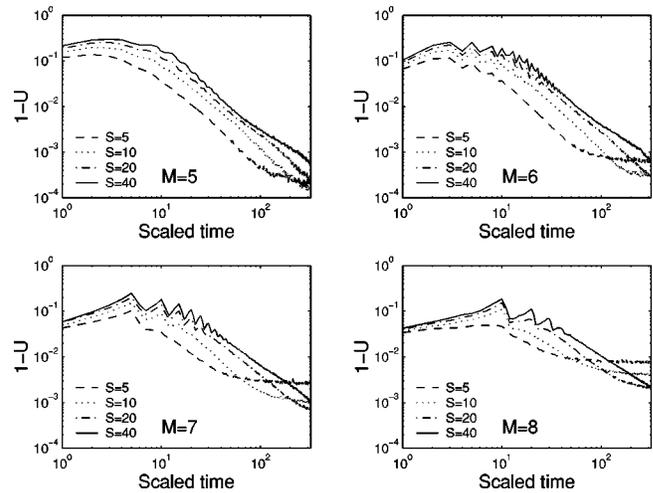


FIG. 6. Utility deviation $1-U$ vs scaled time: Effects of the number of strategies S and memory M . Simulations were done with $N=1001$, $\tau=40$, $n=0.3$, and ensemble averaging each curve over 50 runs.

fit reasonably to a power-law form $1-U \sim (\text{scaled time})^{-x}$, with $x \approx 1.9$. However, in the late time behavior of curves with small adaptation times $\tau=10$ and 20 , we see a slowing down in the utility improvement to apparently constant values of utility, which for $\tau=20$ turns out to be nearer $U_{max}=1$ than for $\tau=10$. On the other hand, for memory $M=7$ and 8 (two lower panels) the overall behavior is fairly similar to that for smaller memory ($M=5$ and 6), though the total improvement in utility appears to be considerably less within the total simulation time used. In fact the change in the efficiency of the games seems to decrease with increasing memory size. This is due to transient period getting longer and utility improvement slowing down or saturating earlier. (In the $M=8$ panel the utility saturation is clearly visible at late times for the adaptation times $\tau=10, 20$, and 40 , while for $\tau=80$ a slowing down in the utility improvement first

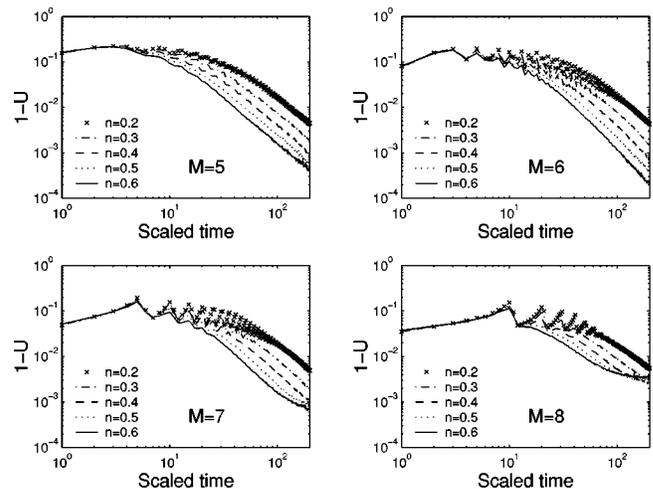


FIG. 7. Utility deviation $1-U$ vs scaled time: Effects of the fraction n of the worst performing agents and memory M . Simulations were done with $N=1001$, $S=10$, $\tau=80$, and ensemble averaging each curve over 50 runs.

appears, possibly turning later to saturation, and for $\tau=160$ is not yet visible but likely to become so even later in scaled time). Between early and late time regions $1-U$ curves show a linear decay with an exponent x comparable to that above, though the fitting is unreliable due to shrinkage of the linear decay region.

It is interesting to note that longer adaptation times eventually lead to a better utility, i.e., a higher efficiency. This is because for a higher dimensional strategy space (due to memory increase) it takes a longer time until sufficiently many histories are gone through to verify the success of a particular strategy. On the other hand, if we do not allow enough time for the adaptation to happen, strategies are changed too often and even the good ones are likely to be disregarded. We will discuss this issue later by exploring the interplay between the adaptation time and simulation or game time.

In Fig. 6 we show the effect of changing the number of strategies $S=5, 10, 20$, and 40 (each panel) and the memory size $M=5, 6, 7$, and 8 (separate panels). We found that for each M value the dependence of the utility deviation, $1-U$ on S , is, in terms of a transient, of linear decay with a fitted power law exponent, and has a late-time saturation behavior quite similar with the dependence on τ , though to a lesser extent. However, unlike previously, increasing the number of strategies seems to shift the utility deviation curve slightly up, such that the system reaches similar efficiency values to those in Fig. 5 later. On the other hand, the dependence of the utility improvement on the memory size is as dramatic as in the previous case.

In Fig. 7 we show the effect of changing the fraction of the worst performing agents $n=0.2, 0.3, 0.4, 0.5$, and 0.6 (each panel) and the memory sizes $M=5, 6, 7$, and 8 (separate panels). Once again we found that the dependence of the utility deviation, $1-U$ on n is in terms of a transient, and of linear decay with a fitted power law exponent, quite similar to the dependence on τ or S . However, increasing the fraction of the worst performing agents seems to push the late-time utility slowing down or saturation to even later scaled times. Once again the effect of memory size is as dramatic as before.

In Figs. 5–7 we see oscillations specifically in panels with larger memory sizes. They seem to be proportional to 2^M , i.e., the number of histories in the game. The reason for these oscillations is most probably the approximately periodic repetition of histories: part of the agents stick to their favorite strategies, which are repeated as a particular history next time around.

Earlier we found that the late-time behavior of the scaled utility changes for different adaptation times. This is because larger memory values increase the dimensionality of the strategy space and thus require longer adaptation times. In order to study this behavior in more detail we have done simulations for several memory sizes $M=5, 6, 7, 8$, and 9 paired either with the total simulation time $T=10\,000$ or $20\,000$. In Fig. 8 we present the results for the utility deviation $1-U$ vs the adaptation time τ , which for a given M shows an overall utility improvement due to doubling the total simulation time. For shorter $T=10\,000$ and all M 's

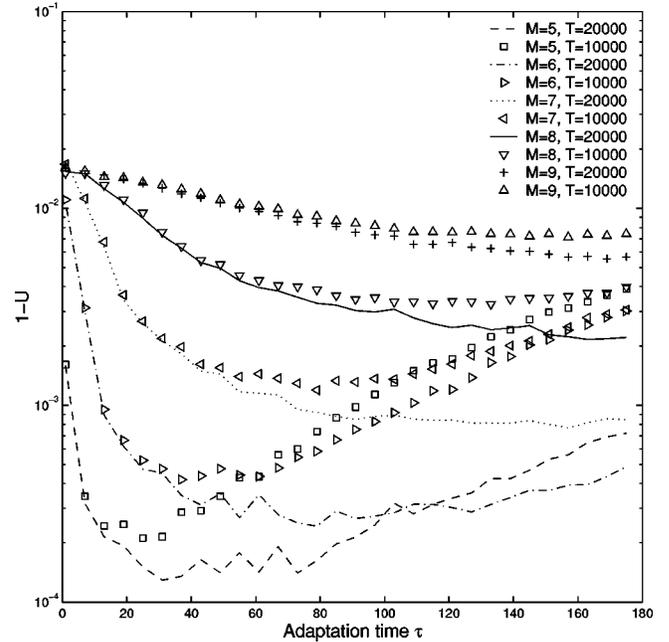


FIG. 8. Utility deviation $1-U$ as a function of the adaptation time τ , for different values of memory M and simulation times (indicated by T). Each point is the time average over the last 500 time steps of the simulation and ensemble averages over 70 runs.

there appear minima in $1-U$ or efficiency maxima, which for $T=20\,000$ and the same M show shifts to higher adaptation times. Thus we expect that increasing the simulation time without limits would make the $1-U$ vs τ curve behave monotonically, approaching an asymptotically maximum efficiency. This can be interpreted that the longer an agent can observe its strategies the more certain it can be of their performances. On the other hand, if the adaptation time is reduced too much, crossovers would take place more at random. For finite simulation times these curves could give guidance for a preferable adaptation time. Intuitively, one could guess that a good adaptation time would be close to 2^M , because if the occurrence of histories were uniformly distributed this would constitute the expectation time for an agent to go through all the histories once and thus see how successful a response determined by a strategy has been in each case.

IV. DISCUSSIONS

In this paper we have studied various genetic algorithm based adaptation mechanisms within the framework of the minority game, and found significant changes in the collective and individual behaviors of the agents. It turned out that the adaptation mechanism in which the best two strategies are chosen as parents and their offspring strategies replace the two worst strategies in the agent's pool leads the system fastest and nearest to maximum utility or efficiency. The pre-eminence of the best adaptation mechanism can be seen at the system level as well as at the agent level: fluctuations in x_t smooth down quickly and the agents outperform those using other mechanisms participating in the same game. The overall success of genetic algorithm based adaptation mecha-

nisms in minority games suggests they also be used in other game theoretic optimization problems. It should be noted that the minority game deviates from the traditional optimization problems because it does not include a particular object function or functions to be maximized. This makes our findings very interesting, exposing a certain characteristic of the minority game, namely, if agents have the possibility to adapt through the responses to the stimuli, they drive towards a state where their own performance improves and the collective of all agents gains a maximum amount of utility every time the game is played. This property stems from convergence of strategies in the strategy space towards, in a way, the optimal ones. They are optimal only in the sense that they tend to bring the maximum utility for the collective,

meaning that at each time step the number of agents winning is as large as possible, i.e., the number of satisfied individual agents is at a maximum. In conclusion it seems possible that simple adaptation schemes and especially those mimicking nature could further extend the class of phenomena which minority game type models would be able to describe.

ACKNOWLEDGMENTS

This research was partially supported by the Academy of Finland, Research Center for Computational Science and Engineering, Project No. 44897 (Finnish Centre of Excellence Programme 2000-2005).

-
- [1] G. Parisi, *Physica A* **263**, 557 (1999).
 [2] B.A. Huberman, P.L.T. Pirolli, J.E. Pitkow, and R.M. Lukose, *Science* **280**, 95 (1998).
 [3] M. Nowak and R. May, *Nature (London)* **359**, 826 (1992).
 [4] T. Lux and M. Marchesi, *Nature (London)* **397**, 498 (1999).
 [5] W.B. Arthur, *Am. Econ. Rev.* **84**, 406 (1994).
 [6] J.H. Holland, *Adaptation in Natural and Artificial Systems* (University of Michigan Press, Ann Arbor, MI, 1975).
 [7] D.E. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning* (Addison-Wesley, Reading, MA, 1989).
 [8] *Handbook of Genetic Algorithms*, edited by D. Lawrence (Van Nostrand Reinhold, New York, 1991).
 [9] D. Challet and Y.-C. Zhang, *Physica A* **246**, 407 (1997).
 [10] D. Challet, M. Marsili, and R. Zecchina, *Phys. Rev. Lett.* **84**, 1824 (2000).
 [11] R. Savit, R. Manuca, and R. Riolo, *Phys. Rev. Lett.* **82**, 2203 (1999).
 [12] A. Cavagna, J.P. Garrahan, I. Giardina, and D. Sherrington, *Phys. Rev. Lett.* **83**, 4429 (1999).
 [13] D. Lamper, S.D. Howison, and N.F. Johnson, *Phys. Rev. Lett.* **88**, 017902 (2002).
 [14] M. Sysi-Aho, A. Chakraborti, and K. Kaski, *Eur. Phys. J. B* **34**, 373 (2003).
 [15] M. Sysi-Aho, A. Chakraborti, and K. Kaski, *Physica A* **322**, 701 (2003); *Phys. Scr.* **T106**, 32 (2003).
 [16] D. Challet and Y.-C. Zhang, *Physica A* **256**, 514 (1998).
 [17] Y. Li, R. Riolo, and R. Savit, *Physica A* **276**, 234 (2000).
 [18] Y. Li, R. Riolo, and R. Savit, *Physica A* **276**, 265 (2000).