

Learning, competition and cooperation in simple games

M.A.R. de Cara^a, O. Pla, and F. Guinea

Instituto de Ciencia de Materiales de Madrid, Consejo Superior de Investigaciones Científicas, 28049 Madrid, Spain

Received 30 June 1999 and Received in final form 27 September 1999

Abstract. The minority model was introduced to study the competition between agents with limited information. It has the remarkable feature that, as the number of strategies available to the agents increases, the collective gain made by the agents is reduced. This crowd effect arises from the fact that only a minority can profit at each moment, while all agents make their choices using the same input. We show that the properties of the model change drastically if the agents make choices based on their individual stories, keeping all remaining rules unaltered. This variation reduces the intrinsic frustration of the model, and improves the tendency towards cooperation and self organization. We finally study the stable mixing of individual and collective behavior.

PACS. 02.50.-r Probability theory, stochastic processes, and statistics – 02.50.Ga Markov processes – 05.40.-a Fluctuation phenomena, random processes, and Brownian motion

The minority game [1] was first introduced in the analysis of decision making by agents with bounded rationality, based on the “*El Farol*” bar problem [2]. A number of agents must make a choice between two alternatives. The choice proves beneficial to a given agent if the total number of agents making that choice is below a given threshold. The game was formulated in a precise way by Challet and Zhang [1]. The bounded rationality of the agents is modeled by assuming that each agent can only process information about the outcomes in the m previous time steps. Given the 2^m possible states an agent could afford, there are 2^{2^m} strategies. Each agent has s strategies, taken at random from the total pool, and for making next decision selects the best performing one of her own set. The choice is successful if the agent is in the minority group, which means that the “comfort” threshold is set at 50% the total number of agents. Finally, the agents assign a score to each strategy at their disposal. The score of the strategies which, at a given time, have predicted the correct outcome is increased by one point.

The game has by now been extensively studied. Particular emphasis has been devoted to the mean square deviation of the number of agents making a given choice, σ , which measures the efficiency of the system. When the fluctuations are large (larger σ), the number of agents in the majority side (the number of losers) increases. In this way, the variance measures the degree of cooperation, or mutual benefit of the agents. It has been shown that it scales with $\rho \equiv 2^m/N$ [1,3–5], where N is the number of agents and 2^m is the number of different configurations that the agents are capable of processing (or *states of the world*, see [6]). When $\rho \gg 1$, the amount of

information available to the agents is so large that they cannot manage to exploit it, and agents take decisions like coin tossing, so that in this limit $\sigma^2/N \rightarrow 1/4$. In the opposite limit, $\rho \ll 1$, the set of strategies of different agents overlap significantly. The agents tend to make similar choices, which puts them often in the majority group. Then σ^2 scales with N^2 , instead of N . This regime is highly inefficient from the point of view of the whole population. The agents manage, however, to arbitrage away all information in the collective history. The value of σ has a minimum for intermediate values of ρ which can be appreciated for not too large values of s . At this minimum, the agents perform better than random, and some degree of cooperation is established. This minimum can be understood as a critical point in an effective spin model with frustrated interactions and an applied field [6].

A crucial ingredient in the model is the fact that all agents act on the same information, irrespective of how it has been generated. Similar results are obtained when the histories are replaced by successions of random numbers [7], which allows for interesting analytical analyses [6,8]. Evolutionary variations, in which agents with different number of strategies, s , capabilities to analyze the time series (as given by m), or additional adjustable parameters have also been studied [3,9]. The $\rho \ll 1$ regime leads not only to large values of σ but also to complex distribution probabilities with a rich structure [10].

The model has been used to describe the interactions of agents competing for scarce resources in different contexts [3,5]. However, it is unlikely that the rules by which the agents make their choices define a evolutionary stable strategy, in the sense commonly used in theoretical

^a e-mail: angeles@quijote.icmm.csic.es

Table 1. Example of strategy $m = 2$.


Signal	Decision
11	1
10	1
01	0
00	1

Table 2. Choices made by an agent playing using the information of the collective history h_c and using the individual history h_i , with the strategy of Table 1. The collective history h_c is the serie of the minority groups, and the individual history h_i is the serie of the choices made by agent i .

	t_{n-2}	t_{n-1}	choice made by agent i at t_n
h_c (minority group)	0	1	0 (Collective MG)
h_i (Group joined by agent i)	1	0	1 (Individual MG)

biology [11]. The low global gain in the limit $\rho \ll 1$ implies that alternative rules can easily improve the performance of the agents. This hypothesis has been verified in different variations of the minority game as defined above. Competition between agents with different memories was first analyzed in [3]. The rules were extended using an additional parameter to improve the chance that the agents use anticorrelated strategies. The value of this parameter was set using an evolution scheme which favors the agent's performance [12]. It has been shown that two populations of agents with different memories, m , perform better than pure populations taken separately [13]. Renewal of the strategies available to the agents also leads to improvements in the performance [14]. In a different context, the global gain made by the agents can increase by adding randomness to the decision making process [8].

We analyze the simplest extension of the model which preserves the basic structure of the agents' decision process. Each agent has the same number of strategies, s , defined in the usual way, which process information from the m preceding time intervals (an example of a strategy when $m = 2$ is given in Tab. 1). The agents, unlike in the usual definition of the game, use as input the time series defined by the individual choices that each agent has made previously. Thus, each agent proceeds differently. An example of the two versions of the minority game is given in Tab. 2. The updating of the score of a given strategy depends on the degree of "comfort" achieved by the agent, that is, on whether the choice made has led it to the minority group. This procedure is identical to that used in the standard version of the model. The effective interaction between agents arises from the updating system in both versions of the game. In the case shown in Table 2, when the minority group is 1, for instance, the strategy shown in table 1 receives one point in the present (individual) model, while it receives 0 points in the standard (collective) minority game.

**Fig. 1.** σ^2/N vs. $2^m/N$ in the collective (three upper graphs) and individual (three lower ones) games. Each point represents the average of 5 independent runs for different values of N , and $m = 4$ (\triangle), $m = 5$ (\blacktriangle), $m = 6$ (\circ), $m = 7$ (\bullet), $m = 8$ (\square), $m = 9$ (\blacksquare). For clarity each value of s is represented in a separate graph.

We think that the extension of the minority game presented here is relevant to many situations, assuming a context of incomplete information, or bounded rationality. In the absence of an exact procedure to optimize individual choices, we find very plausible that agents will consider only their previous actions, or those of a small subgroup of all the players, close to them by some kind of affinity [16]. Obviously, if the outcome is determined by the global behavior, the tendency to repeat those actions which lead to the greatest comfort will be influenced by the dynamics of the whole population.

While the extension of the minority game presented here is relatively straightforward, it changes significantly some of the features needed for analytical models to work. As discussed below, our model does not satisfy the independence of histories found in the standard minority game [7] (see below). Hence, the mapping of the problem into an effective spin glass model [6] is not possible. As discussed in detail, our numerical results show no evidence of a sharp phase transition in the range of parameters where it has been found for the standard minority game. Using the language of statistical mechanics, our model belongs to a different universality class. Our results suggest that the most relevant feature of a spin glass, the presence of frustration, is greatly suppressed. Thus, our version of the minority game provides a simple scheme in which a collective of selfish agents can avoid, to some extent, the “tragedy of the commons” [15].

In order to characterize the various regimes of our model, we compare the values of the mean square deviations of the attendances, σ , in the present version of the minority game and the canonical results in Figure 1.

In the limit when the information available to the agents is too large, we find that $\sigma^2 \rightarrow N/4$, the same result as if the agents made their choices at random, just in the case of high values of s . For case $s = 3$ shown in the figure $\sigma^2 \rightarrow N/5$, and this value is even lower for $s = 2$. $s = 1$ is, as in the standard game, highly sensible to the initial conditions, and averaging over them, gives a dispersion equal to $N/4$ independently of ρ . In the limit $\rho \rightarrow 0$, the values of σ are significantly lower in the “individual” version of the game presented here, and comparable, or lower, than those found in other extensions of the model. There is a significant spreading as function of m and N , meaning that the scaling with ρ is not too well satisfied. The scaling with ρ implicitly assumes that all possible histories appear with the same probability in the collective history [4,6,7]. In the present version of the model, if the individual stories used by the agents are replaced by random series, σ takes values close to the random case, irrespective of the value of ρ . Thus, the main hypothesis used to justify the scaling in the minority game in its usual form does not hold in this case.

The group which was on the winning side can be inferred from the “comfort” that the agent gained after each outcome. This information is used in updating the score of the strategies, which, however, act on a different input. As this input is not the same for all agents, they have no obstacle in following anticorrelated dynamics, even when all use similar strategies. The measure of that correlation can be analyzed explicitly by taking the average Hamming distance between agents histories [16]. We have further analyzed this point by calculating the average number of histories processed by the agents. The number of histories is always significantly below that in the canonical model ($P = 2^m$), implying that the system tends to be locked into situations where agents generate a relatively small number of possibly anticorrelated individual histories. This P , is also a function of m , N , and s , in such a way that it decreases monotonically when in-




Fig. 2. Mean square deviation of the attendances in a model where agents use collective and individual rules (left axis), and fraction of agents which use an individual rule (right axis). Different symbols correspond to different choices of m (see Fig. 1).

creasing N and decreasing m . When s is small the limit for large m and small N , is not 2^m , but some lower value. This would explain the limit of $\sigma^2/N \neq 1/4$ when ρ is large discussed above.

The present version of the model needs not to define an evolutionary stable strategy. If there is information available in the series of global minority groups, an agent playing according to the canonical rules will benefit from doing so. We have analyzed the competition between these two types of behavior by allowing each agent to have a dual scoring system for its strategies, following the two set of rules. Each agent plays the strategy with the highest score at a given time step. Thus, the population can be divided into those using collective rules and those using individual rules. The values of σ obtained in this way, and the fraction of agents using a collective strategy are shown in Figure 2.

In the limit when the information available is large, $\rho \rightarrow \infty$, we recover the random value for σ . Then, both behaviors are indifferent, and the agents use 50% of the time each of them. The fraction of agents which use a collective behavior has a maximum near the value of ρ for which σ has a minimum in the usual version




Fig. 3. The same magnitudes drawn in 2, but for the case $s = 1$. Also different symbols correspond to different choices of m (see Fig 1).

of the model. Finally, the number of agents using collective rules strongly decreases as $\rho \rightarrow 0$. In this limit, the preferable behavior is the individual one outlined here, although a small fraction of agents using a collective approach survives. The global efficiency, however, is decreased. Thus, although a mixed population is the stable situation, the small fraction of agents which follow collective rules behave in a parasitic way, lowering the overall gain.

The most striking difference with the usual version of the minority game takes place when only one strategy is available to each agent, $s = 1$. This case is trivial in the minority game, as the agents have no way to learn or to adapt. The same applies if each agent uses a purely individual set of rules. When the agents can use the best of the two behaviors, the strategy of each agent can be used to process two inputs: the collective history of winning sides, or the succession of prior choices made by that agent. This is shown in Figure 3.

The global performance of an hybrid set of agents using both collective and individual rules is best when $s = 1$ for a large range of values of ρ . A qualitative explanation of the adaptability of the agents in this extreme limit can be obtained by noting that, when a given agent repeatedly makes an incorrect choice, its individual history is anticorrelated with the sequence of collective best choices. Thus, if the strategy at its disposal gives a different out-

come when presented with the two inputs, the agent will tend to give the opposite answer to that used, unsuccessfully, before. There is a self correcting mechanism built into the model, which tends to prevent very negative performances. On the other hand, if the agents are locked into a situation where each of them obtains about 50% of the points, a stable situation can be achieved, where the agents remain anticorrelated by alternating between the two inputs at the disposal of each of them. This is consistent with the result that the fraction of agents using collective and individual behavior is comparable for all values of ρ .

In conclusion, we have discussed the simplest extension of the minority game which preserves the basic parameters of the model. We show that agents with the same processing power as in the usual model can perform much better if they use their individual histories as input, instead of the evolution of the global system. An evolutionary stable situation arises with agents which can use both collective and individual rules. The capability of the agents to adapt and increase the global performance is significantly enhanced, and herd effects disappear. These emergent features change qualitatively even the simplest and most trivial version of the minority game, that in which each agent disposes of a single strategy.

Financial support from the European Union through grant ERB4061PL970910 and the Spanish Administration through grant PB97/0875 are gratefully acknowledged. We acknowledge useful conversations with D. Challet, M. Marsili and Yi-Cheng Zhang.

References

1. D. Challet, Y.-C. Zhang, *Physica A* **246**, 407 (1997).
2. W.B. Arthur, *Am. Econ. Assoc. Papers Proc.* **84**, 406 (1994).
3. D. Challet, Y.C. Zhang, *Physica A* **256**, 514 (1998).
4. Y.-C. Zhang, *Europhysics News* **28**, 51 (1998).
5. R. Savit, R. Manuca, R. Riolo, *Phys. Rev. Lett.* **82**, 2203 (1999).
6. D. Challet, M. Marsili, *Phys. Rev. E*, **60**, R6271 (1999); D. Challet, M. Marsili, R. Zechina, preprint (cond-mat/9904392).
7. A. Cavagna, *Phys. Rev. E*, **59** R3783 (1999).
8. A. Cavagna, J.P. Garrahan, I. Giardina, D. Sherrington, *Phys. Rev. Lett.* **83**, 4429 (1999).
9. N.F. Johnson, S. Jarvis, R. Jonson, P. Cheung, Y.R. Kwong, P.M. Hui, *Physica A*, **256**, 230 (1998). N.F. Johnson, M. Hart, P.M. Hui, *Physica A*, **269**, 1 (1999).
10. M.A.R. de Cara, O. Pla, F. Guinea, *Eur. Phys. J. B*, **10**, 187 (1999).
11. J. Maynard Smith, *Evolution and the theory of games* (Cambridge U. P., Cambridge, 1982).
12. N.F. Johnson, P.M. Hui, R. Jonson, T.S. Lo, *Phys. Rev. Lett.* **82**, 3360 (1999).
13. N.F. Johnson, P.M. Hui, D. Zheng, M. Hart, *J. Phys. A* **32**, L427 (1999).
14. Y. Li, R. Riolo, R. Savit, preprint (cond-mat/9903008).
15. G. Hardin, *Science* **162**, 1243 (1968).
16. M.A.R. de Cara, O. Pla, F. Guinea (forthcoming).