

Self-Segregation versus Clustering in the Evolutionary Minority Game

Shahar Hod¹ and Ehud Nakar²

¹*Department of Condensed Matter Physics, Weizmann Institute, Rehovot 76100, Israel*

²*The Racah Institute of Physics, The Hebrew University, Jerusalem 91904, Israel*

(Received 26 December 2001; published 23 May 2002)

Complex adaptive systems have been the subject of much recent attention. It is by now well established that members (“agents”) tend to self-segregate into opposing groups characterized by extreme behavior. However, the study of such adaptive systems has mostly been restricted to simple situations in which the prize-to-fine ratio R equals unity. In this Letter we explore the dynamics of evolving populations with various different values of the ratio R , and demonstrate that extreme behavior is in fact *not* a generic feature of adaptive systems. In particular, we show that “confusion” and “indecisiveness” take over in times of depression, in which case cautious agents perform better than extreme ones.

DOI: 10.1103/PhysRevLett.88.238702

PACS numbers: 02.50.Le, 87.23.Kg, 89.65.Ef

A problem of central importance in social, biological, and economic sciences is that of an evolving population in which individual agents adapt their behavior according to past experience, without direct interaction between different members. Of particular interest are situations in which members (usually referred to as “agents”) compete for a limited resource, or to be in a minority (see, e.g., [1], and references therein.) In financial markets, for instance, more buyers than sellers implies higher prices, and it is therefore better for a trader to be in a minority group of sellers. Predators foraging for food will do better if they hunt in areas with fewer competitors. Rush-hour drivers, facing the choice between two alternative routes, wish to choose the route containing the minority of traffic [2].

Considerable progress in the theoretical understanding of such systems has been gained by studying the simple, yet realistic model of the minority game (MG) [3], and its evolutionary version (EMG) [1] (see also [4–13], and references therein). The EMG consists of an odd number of N agents repeatedly choosing whether to be in room “0” (e.g., choosing to sell an asset or taking route A) or in room “1” (e.g., choosing to buy an asset or taking route B). At the end of each turn, agents belonging to the smaller group (the minority) are the winners, each of them gains 1 point (the “prize”), while agents belonging to the majority room lose 1 point (the “fine”). The agents have a common “memory” look-up table, containing the outcomes of m recent occurrences (the particular value of m is of no importance [1]). Faced with a given bit string of recent m occurrences, each agent chooses the outcome in the memory with probability p , known as the agent’s “gene” value (and the opposite alternative with probability $1 - p$). If an agent score falls below some value d , then its strategy (i.e., its gene value) is modified. In other words, each agent tries to learn from his past mistakes, and to adjust his strategy in order to survive.

A remarkable conclusion deduced from the EMG [1] is that a population of competing agents tends to self-segregate into opposing groups characterized by extreme behavior. It was realized that in order to flourish in such

situations an agent should behave in an extreme way ($p = 0$ or $p = 1$) [1].

It should be emphasized, however, that previous analyses were restricted to the simple case in which the prize-to-fine ratio R was assumed to be equal unity. On the other hand, in many real life situations this ratio may take a variety of different values. In the extreme situation, the fine (e.g., a temporary worker getting fired of work after being late to the office due to a traffic jam, or a predator being starved to death while unsuccessfully trying to hunt in an area with many competitors) may be larger than the prize (a day’s payment or a successful hunt which guarantees food for few days, respectively). Another example is that of a trader in a financial market which is under depression. In such circumstances, the trader usually loses more money in a bad deal than he gains in a successful one (due to overall reduction in market’s value).

Moreover, we know from real life situations that extreme agents not always perform better than cautious ones. In particular, our daily experience indicates that in difficult situations (e.g., when the prize-to-fine ratio is low) human people tend to be confused and indecisive. In fact, in such circumstances they usually seek to do the *same* (rather than the opposite) as the majority.

Thus, of great interest for real social and biological systems are situations in which the prize-to-fine ratio is smaller (or larger) than unity. The aim of the present Letter is to explore the dynamics of evolving populations with various different external conditions (i.e., different values of the ratio R). Of main importance is the identification of the strategies that perform best in a particular situation.

Figure 1 displays the long-time frequency distribution $P(p)$ of the agents [the lifespan, $L(p)$, defined as the average length of time a strategy p survives between modifications, has a similar behavior]. We find three qualitatively different populations, depending on the precise value of the prize-to-fine ratio parameter R . For $R > R_c^{(1)}$ (this includes the case studied so far in the literature, $R = 1$. The value of $R_c^{(1)}$ depends on the number of agents and the parameter d) the distribution becomes peaked around $p = 0$

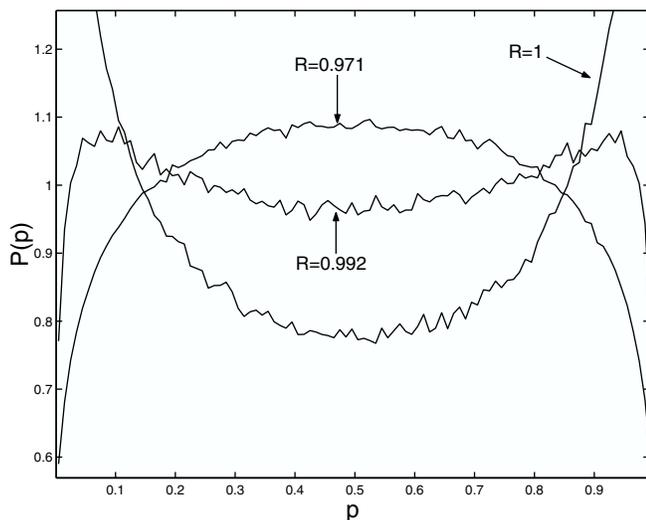


FIG. 1. The strategy distribution $P(p)$ for different values of the prize-to-fine ratio: $R = 0.971$, $R = 0.992$, and $R = 1$. The results are for $N = 10001$ agents, $d = -4$. Each point represents an average value over 10 runs and 100 000 time steps per run.

and $p = 1$; the population will self-segregate (this corresponds to always or never following what happened last time). To flourish in such a population, an agent should behave in an *extreme* way [1]. On the other hand, for $R < R_c^{(2)}$ (poor conditions, in which the fine is larger than the reward) the population tends to crowd around $p = \frac{1}{2}$. This corresponds to “confused” and “indecisive” agents. There is also an intermediate phase [for $R_c^{(2)} < R < R_c^{(1)}$], in which the population tends to form an M-shaped distribution, peaked around some finite p_0 and its counterpart $1 - p_0$ (with the absolute minimas of the distribution located at $p = 0$ and $p = 1$).

An important feature of the original EMG (for the $R = 1$ case [1]) is that the root-mean-square (rms) separation of the strategies is *higher* than the corresponding value for uniform $P(p)$. This indicates the desire of agents to do the *opposite* of the majority [1]. Figure 2 shows the rms separation of the population as a function of the prize-to-fine ratio R . Remarkably, we find that for small values of R the rms is in fact *smaller* than that obtained for a uniform $P(p)$ distribution. We therefore conclude that in times of difficulties agents desire to do the *same* (rather than the opposite) as the majority. This is exactly the type of behavior we anticipated in the introduction based on daily experience.

Qualitatively, we have found that the larger the number of agents N , the sharper the dependence of the system’s behavior is on the prize-to-fine ratio R (that is, the peak in Fig. 2 is sharper for larger values of N). This may indicate a sharp phase transition (instead of a continuous change of the global behavior with R) in the limit of $N \rightarrow \infty$.

In the original EMG [1] it was found that the dynamics of the system leads to situations in which the size of the minority group is maximized, indicating that the efficiency of the system is maximized. The (scaled) efficiency of the

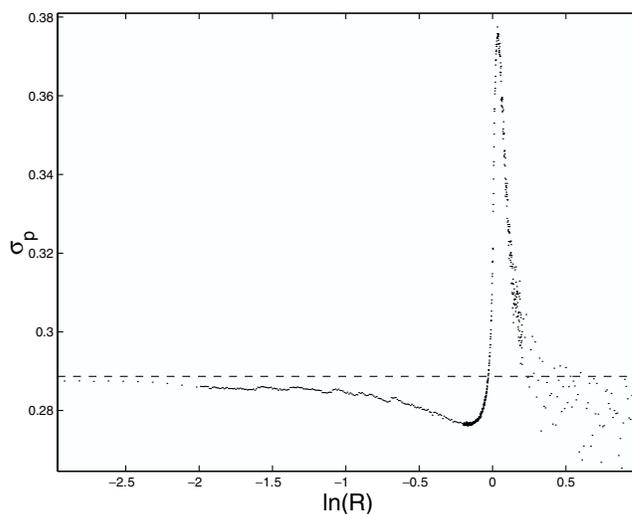


FIG. 2. The root-mean-square separation (rms) of the strategies as a function of the prize-to-fine ratio R . The horizontal line represents the rms separation for a uniform $P(p)$ distribution. $N = 1001$, $d = -4$. Each point represents an average value over 10 runs and 10 000 time steps per run.

system is defined as the number of agents in the minority room, divided by the maximal possible size of the minority group, $(N - 1)/2$. Figure 3 displays the system’s efficiency as a function of the ratio R . We also display the efficiency for agents guessing *randomly* between room 0 and room 1, and for a *uniform* distribution of agents. As previously found, there is a range of R (which includes the previously studied case $R = 1$ [1]) in which the efficiency of the system is *better* than the random case. However, for small values of the prize-to-fine ratio, the efficiency of the system is remarkably *lower* than that obtained for agents choosing via independent coin tosses. Thus, considering the efficiency of the system as a whole, the agents would

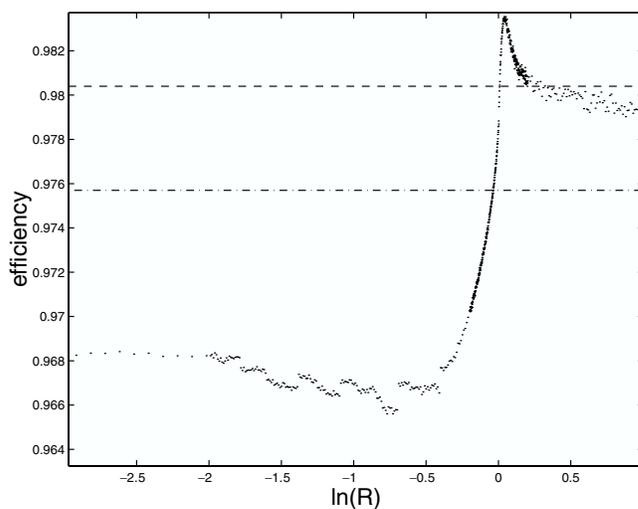


FIG. 3. The (scaled) efficiency E of the system as a function of the prize-to-fine ratio R . Horizontal lines represent the efficiency for uniform $P(p)$ distribution (dashed) and a coin-tossing situation (dash-dotted). The parameters are the same as in Fig. 2.

be better off not adapting their strategies because they are doing *worse* than just guessing at random.

Note that an optimum utilization of the resources is obtained at some $R_{\max} > 1$ (with $R_{\max} - 1 \ll 1$). This implies that an evolving population requires a small positive feedback in order to exploit its resources in an optimal way. On the other hand, a wealthy society has an efficiency which is worse than that of a uniform $P(p)$ distribution (this occurs for prize-to-fine ratios which are too large). This reflects the fact that in a “spoiling” environment the agents have no real motivation to evolve (they have a long lifespan even without exploiting their resources in an optimal way).

In previous studies (of the $R = 1$ case) it has been established that the evolving population enters into a *stationary* phase, in which case the $P(p)$ distribution remains essentially constant in time [4,7]. In Fig. 4 we display the time dependence of the average gene value, $\langle p \rangle$, for different values of the prize-to-fine ratio R . The distribution $P(p)$ oscillates around $p = \frac{1}{2}$. The smaller the value of R the larger are the amplitude and the frequency of the oscillations. Thus, we conclude that a population which evolves in a tough environment never establishes a steady state distribution. Agents are constantly changing their strategies, trying to survive. By doing so they create global currents in the gene space.

We now provide some analytical analysis of the problem, a generalization of the one presented in [1] for arbitrary values of the prize-to-fine ratio R . The simplest example of our system contains $N = 3$ agents, and three discrete gene values $p = 0, \frac{1}{2}, 1$. We consider configurations for which the average gene value lies between $\frac{1}{3}$ and $\frac{2}{3}$, a reminiscent of the fact that $\langle p \rangle$ displays only mild oscillations around $p = \frac{1}{2}$. To obtain the average $P(p)$ distribution we weigh the various configurations according to the average points awarded per agent in each of

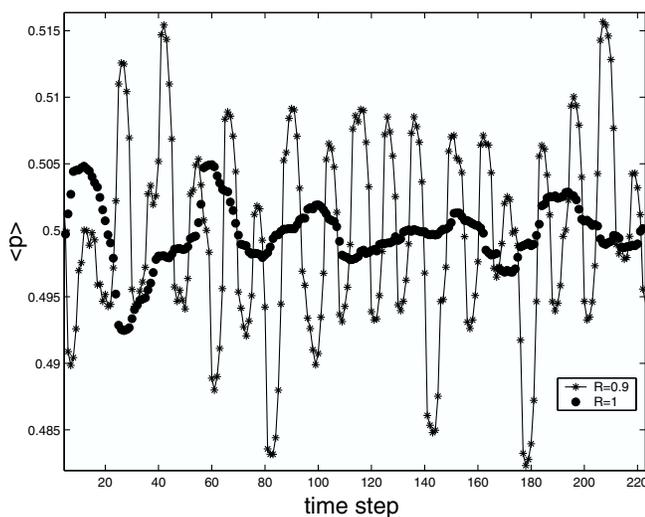


FIG. 4. Time evolution of the average gene value $\langle p \rangle$, for different values of the prize-to-fine ratio: $R = 0.9$ and $R = 1$. The parameters are the same as in Fig. 1.

the configurations [14]. The analytical results are given in Table I [note that $P(0) = P(1)$]. We find that this simplified toy model provides a fairly good qualitative description of the complex system. In particular, it follows that the population self-segregates for prize-to-fine ratios larger than $R_c = \frac{5}{7}$, while for $R < R_c$ the agents tend to cluster around $p = \frac{1}{2}$. In addition, the efficiency of the system is maximal at *intermediate* values of the prize-to-fine ratio, while poor ($R < 2$) and wealthy ($R > 3$) populations display a lower efficiency.

One can improve the analysis of the evolutionary minority game with the aid of a semianalytical model [15]. The semianalytical model is based on the fact that the population never establishes a true stationary distribution (see Fig. 4). This fact has been ignored in previous studies of the evolutionary minority game. Thus, the probability of a particular agent to win is time dependent. In fact, the winning probability oscillates with time; the oscillation amplitude depends on both the value of the prize-to-fine ratio R , and on the agent’s gene value p (the smaller the value of R , the larger is the oscillation amplitude; in addition, agents with $p = 0, 1$ have an oscillation amplitude larger than those with $p = 1/2$).

It turns out that the temporal oscillations of the winning probabilities explain the transition of the system’s global behavior from self-segregation to clustering: For small values of R (when the prize is smaller than the fine) there is a negative global drift of the agents’s score towards $-d$. Thus, agents with an equal number of winnings and losses will eventually perish [after approximately $2d/(1 - R)$ turns]. In order to survive in harsh conditions ($R < 1$) agents must win more times than they lose. Agents with $p = 0, 1$ have a winning probability which oscillates in time with a large amplitude, and therefore most of these agents win and lose the same number of times in each cycle of the oscillations. On the other hand, agents with $p = 1/2$ have a winning probability which is practically constant ($\approx 1/2$) in time. Thus, these are the only agents that can win more times than they lose. Therefore, for small values of the prize-to-fine ratio, agents tend to cluster around $p = 1/2$.

In summary, we have explored the evolution of complex adaptive systems with an arbitrary value of the prize-to-fine ratio R . The main results and their implications are as follows:

(i) It has been widely accepted that *self-segregation* is a generic characteristic of an evolving population of competing agents. This belief was based on studies of the

TABLE I. Distribution of strategies and efficiency of a three agents system.

R	$P(0):P(\frac{1}{2})$	Efficiency
$R < 2$	$\frac{19R-53}{26R-58}$	$\frac{57R-150}{64R-164}$
$2 \leq R \leq 3$	2.5	1
$R > 3$	$\frac{23R-49}{23R-61}$	$\frac{85R-191}{92R-212}$

$R = 1$ case. Our analysis, however, turns over this point of view. In particular, in times of difficulties agents tend to *cluster* around $p = \frac{1}{2}$; cautious agents perform better (live longer) than extreme ones. Stated in a more pictorial way, confusion and indecisiveness take over at tough times.

(ii) In previous analyses it was found that agents desire to do the opposite of the majority [1]. We have shown that this property is in fact not a generic one. In particular, in a tough environment agents try to do the *same* as the majority [the rms separation of strategies is in fact *smaller* than that obtained for a uniform $P(p)$ distribution].

(iii) For small values of the prize-to-fine ratio (poor external conditions) the efficiency of the system is well below the efficiency achieved by random agents (ones who choose via independent coin tosses). It seems that “panic” and “confusion” (clustering around $p = \frac{1}{2}$) prevent the agents from achieving a reasonable utilization of resources. Similarly, a wealthy population, for which there is no real motivation for adaptation, displays a poor efficiency. On the other hand, an evolving population achieves an optimum utilization of its resources when it receives a (*small*) positive external reinforcement (that is, for $0 < R_{\max} - 1 \ll 1$).

(iv) The gene distribution $P(p)$ displays temporal oscillations around $p = \frac{1}{2}$. The smaller the value of the prize-to-fine ratio, the farther the system is from a steady-state distribution. This in particular implies that the steady-state assumption used to analyze the EMG (in the $R = 1$ case) [7] is no longer valid for smaller values of R .

S. H. thanks Mordehai Milgrom for his kind assistance, and a support by the Dr. Robert G. Picard fund in physics. The research of S.H. was supported by Grant No. 159/99-3 from the Israel Science Foundation.

- [1] N.F. Johnson, P.M. Hui, R. Jonson, and T.S. Lo, Phys. Rev. Lett. **82**, 3360 (1999).
- [2] B. Huberman and R. Lukose, Science **277**, 535 (1997).
- [3] D. Challet and C. Zhang, Physica (Amsterdam) **246A**, 407 (1997); **256A**, 514 (1998); **269A**, 30 (1999).
- [4] R. D’Hulst and G. J. Rodgers, Physica (Amsterdam) **270A**, 514 (1999).
- [5] H. Ceva, Physica (Amsterdam) **277A**, 496 (2000).
- [6] E. Burgos and H. Ceva, Physica (Amsterdam) **284A**, 489 (2000).
- [7] T. S. Lo, P.M. Hui, and N.F. Johnson, Phys. Rev. E **62**, 4393 (2000).
- [8] P.M. Hui, T. S. Lo, and N.F. Johnson, e-print cond-mat/0003309.
- [9] M. Hart, P. Jefferies, N.F. Johnson, and P.M. Hui, e-print cond-mat/0003486; e-print cond-mat/0004063.
- [10] E. Burgos, H. Ceva, and R.P.J. Perazzo, e-print cond-mat/0007010.
- [11] T. S. Lo, S. W. Lim, P.M. Hui, and N.F. Johnson, Physica (Amsterdam) **287A**, 313 (2000).
- [12] Y. Li, A. VanDeemen, and R. Savit, e-print nlin.AO/0002004.
- [13] R. Savit, R. Manuca, and R. Riolo, Phys. Rev. Lett. **82**, 2203 (1999).
- [14] For $R < 2$ all configurations have a negative average yield. The relative weight of each configuration is inversely proportional to its yield. For $R \geq 2$ we weight those configurations that have a positive average gain, the weight being proportional to this yield. Consider, for instance, the configuration in which there is one agent in each of the gene values. In this case, each of the extreme agents wins with probability $\frac{1}{2}$, whereas the “confused” agent ($p = \frac{1}{2}$) cannot win. Thus, the average number of points gained per agent per time step is $(R - 2)/3$. The relative weight of this configuration is therefore proportional to $[(R - 2)/3]^{-1}$ in the $R < 2$ case, and to $(R - 2)/3$ for $R \geq 2$.
- [15] S. Hod and E. Nakar (to be published).