Adaptive Competition, Market Efficiency, and Phase Transitions

Robert Savit, Radu Manuca, and Rick Riolo
Program for the Study of Complex Systems and Physics Department, University of Michigan, Ann Arbor, Michigan 48109
(Received 3 August 1998)

In many social and biological systems agents simultaneously and adaptively compete for limited resources, thereby altering their environment. We analyze a simple model that incorporates fundamental features of such systems. If the space of strategies available to the agents is small, the system is in a phase in which all information available to the agents' strategies is traded away, and agents' choices are maladaptive, resulting in a poor collective utilization of resources. For larger strategy spaces, the system is in a phase in which the agents are able to coordinate their actions to achieve a better utilization of resources. The best utilization of resources occurs at a critical point, when the dimension of the strategy space is on the order of the number of agents. [S0031-9007(99)08619-6]

PACS numbers: 02.50.Le, 05.65.+b, 05.70.Fh, 87.23.Ge

Most systems in the biological and social sciences involve interacting agents, each making behavioral choices in the context of an environment that is formed, in large part, by the collective action of the agents themselves, and with no centralized controller acting to coordinate agent behavior. In the most interesting cases, the agents have heterogeneous strategies, expectations, and beliefs [1]. In some cases, the agents' strategies may be self-validating, at least for a limited time. For example, in the financial markets a widespread belief that a commodity will rise in price may perforce result in a price rise for that commodity. But unless there are fundamental reasons for the price rise, such bubbles eventually burst, so that widely shared strategies are often self-defeating in the long run. Thus, in many systems successful agents will employ strategies that differentiate them from their competitors. Furthermore, from the point of view of overall system performance, the best strategy sets are those that result in coordinated resource utilization so that average agent experience is relatively good, and resources are consumed near their limiting rates. Examples of systems in which agents seek to differentiate themselves from their competitors, and in which coordinated allocation of resources is critical, include firms searching for profitable technological innovations, ecological communities [2], routers sending packets over the internet [3], and humans deciding on which night to go to a popular bar [1].

Although these systems are enormously complicated, there are fundamental properties which are shared by all of them. To understand such systems, we must first understand the dynamics imposed by their most basic common properties.

The simple model of competition we discuss here [4] consists of N agents playing a game as follows: At each time step of the game, each of the N agents joins one of two groups, labeled 0 or 1. Each agent that is in the minority group at that time step is awarded a point, while each agent belonging to the majority group gets nothing. An agent chooses which group to join at a given time step based on the prediction of a strategy. The strategy uses information from the historical record of which group was the minority group as a function of time. A strategy of memory m is a table of two columns and 2^m rows. The left column contains all of the 2^m possible combination of m 0's and 1's, while each entry in the right column is a 0 or a 1. To use this strategy, an agent observes which groups were the minority groups during the immediately preceding m time steps, and finds that string of 0's and 1's in the left column of the table. The corresponding entry in the right column contains that strategy's determination of which group (0 or 1) the agent should join during the current time step.

In each of the games discussed here, all strategies used by all of the agents have the same value of m. At the beginning of the game each agent is randomly assigned s (>1) of the 2^{2^m} possible strategies, chosen with replacement [5]. For its current play the agent chooses its strategy that would have had the best performance over the history of the game until that time. Ties between strategies are decided by a coin toss. Because the agents each have more than one strategy, the game is adaptive in that agents can choose to play different strategies at different moments of the game in response to changes in their environment. Because the environment (i.e., the time series of minority groups) is created by the collective action of the agents themselves, and because the relative rankings of the agents' strategies depend on their previous successes, this system has strong collective feedback.

This system may be though of as a very simple model for a number of different situations in the social and biological sciences. In particular, this system can be interpreted as a kind of very simple "protomarket," driven by a simple supply-demand dynamic [6].

We report here the results of this game for a range of values of N (odd), m, and $s = 2$. The qualitative results also hold for other values of s that are not extremely large [7]. We must also create a short (of order m) random history of 0's and 1's, so that the strategies can be initially evaluated. The asymptotic statistical results of any run do not materially depend on what this random string is.

To understand the behavior of this system, consider the time series of the number of agents belonging to group 1, which we will call L_1. (This information is not available to the agents but it is available to the researchers.) The mean of this series is generally close to $N/2$ for all values of N, m, and s, so the standard deviation σ of this time series is a convenient reciprocal measure of how effective the system is at distributing resources, on average, since the smaller σ is, the larger a typical minority group is.

The behavior of σ is quite remarkable. In Fig. 1, we plot σ for these time series as a function of m for $N = 101$ and $s = 2$. For each value of m, 32 independent runs were performed. The horizontal dashed line in this graph is at the value of σ for the random choice game (RCG), i.e., for the game in which each agent randomly chooses 0 or 1, independently and with equal probability at each time step. Note the following features: (1) For small m, the average value of σ is very large (much larger than in the RCG). In addition, for $m < 6$ there is a large spread in the σ's for different runs with different (random) initial distributions of strategies to the agents, but with the same m. (2) There is a minimum in σ at $m = 6$ at which σ is less than the standard deviation of the RCG. We shall refer to the value of m at which the σ vs m curve (for fixed N) has its minimum as m_c. Also, for $m \geq m_c$, the spread in the σ's appears to be small relative to the spread for $m < m_c$. (3) As m increases beyond 6, σ slowly increases, and for large m approaches the value for the RCG.

The system clearly behaves in a qualitatively different way for small and large m. To understand the dynamics in these two regions, consider the (binary) time series of minority groups G, the data publicly available to the agents. To study the information content of G, consider $P(1 \mid u_k)$, the conditional probability to have a 1 immediately following some specific string, u, of k elements. Recall that in a game played with memory m, the strategies use only the information encoded in strings of length m to make their choices. In Fig. 2, we plot $P(1 \mid u_k)$ for G generated by a game with $m = 4$, $N = 101$, and $s = 2$. Figure 2(a) shows the histogram for $k = m = 4$ and Fig. 2(b) shows the histogram for $k = 5$. Note that the histogram is quite flat at 0.5 in Fig. 2(a), but is not flat in Fig. 2(b). Thus, for any strategy with memory (less than or) equal to 4, the history of minority groups contains no predictive information about which will be the minority group at the next time step. But recall that this time series itself was generated by players playing strategies with $m = 4$. Therefore, in this sense, the market is efficient [8] and no strategy using memory (less than or) equal to 4 can, over the long run, have a success rate better than 50%. But G is not a random (IID) sequence. There is information in G, as indicated by the fact that the histogram in Fig. 2(b) is not flat. However, that information is not available to the strategies of the agents playing the $m = 4$ game who collectively generated G in the first place. We shall refer to this property as "strategy efficient" to distinguish it from other kinds of market efficiency [8].

We can repeat this analysis for $m \geq 6$ ($N = 101$, $s = 2$). For this range of m, the corresponding histogram for $k = m$ is not flat, as we see in Fig. 3 for the $m = 6$ game. In this case, there is significant information available to the strategies of the agents playing the game with memory m and the market is not efficient in this sense.

How does the system behavior depend on N? One finds, plotting σ vs m for each fixed N, that in all cases one obtains a graph with a shape similar to that of Fig. 1, but

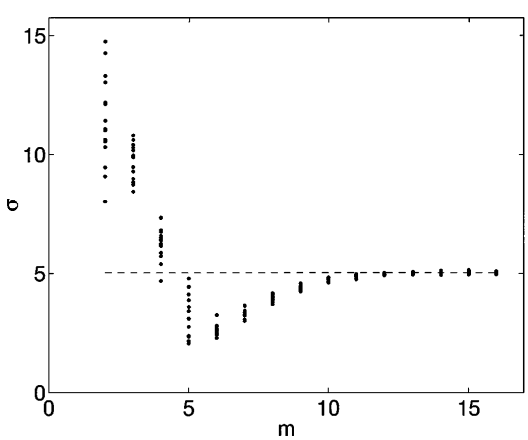

FIG. 1. σ as a function of m for $N = 101$ and $s = 2$, showing 32 independent runs of 10 000 time steps for each value of m. The value of σ for each run is indicated by a dot. The horizontal dashed line is at the value of σ for the random game described in the text.

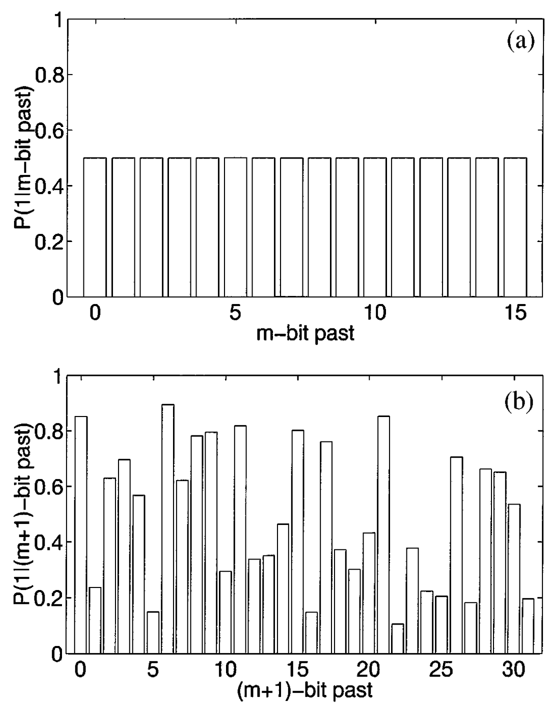

FIG. 2. (a) A histogram of the conditional probability $P(1 \mid u_k)$ with $k = 4$ for the game played with $m = 4$. The bin numbers, when written in binary form, yield the strings u. (b) A histogram of the conditional probability $P(1 \mid u_k)$ with $k = 5$ for the game played with $m = 4$.

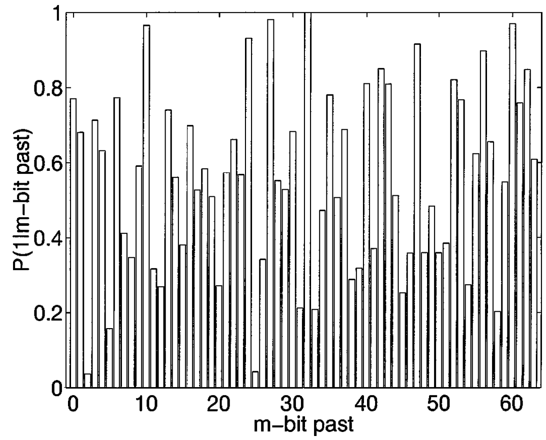

FIG. 3. A histogram of the conditional probability $P(1 \mid u_k)$ with $k = 6$ for the game played with $m = 6$. The bin numbers, when written in binary form, yield the strings u.

in which the position of the minimum, m_c, is proportional to $\ln N$. In addition, σ and the spread in σ behave in very simple ways as a function of N. Generally, for fixed $m < m_c$, both σ and the standard deviation of the σ's (defined as $\Delta\sigma$) are proportional to N, whereas for fixed $m \geq m_c$ both σ and $\Delta\sigma$ are proportional to $N^{1/2}$.

The transition between these very different behaviors is at $m_c \sim \ln N$. We have found [6], using mean-field-like arguments, that to a first approximation σ^2/N is a function only of $2^m/N \equiv z$. To see this explicitly, we plot in Fig. 4 σ^2/N as a function of z on a log-log scale for various N and m (with $s = 2$). Note that all of the data fall on a nearly universal curve. The minimum of this curve is near $2^{m_c}/N \equiv z_c \approx 0.5$, and separates the two different phases. The slope for $z < z_c$ approaches -1 for small z, while the slope for $z > z_c$ approaches zero for large z, consistent with the results of Fig. 4 [9]. Because σ^2/N depends only on z, it is clear that, for any fixed z, σ is proportional to $N^{1/2}$, both above and below z_c. In addition, it can be shown that, for fixed z, $\Delta\sigma$ is approximately independent of N, approaching a z-dependent constant as $N \rightarrow \infty$. In the $N \rightarrow \infty$ limit, $\Delta\sigma$ is large for small values of z and decreases monotonically with increasing z. It is unclear whether or not $\Delta\sigma$ is nonanalytic at z_c.

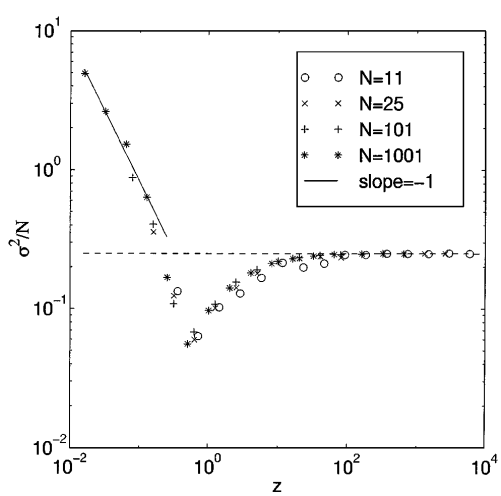

FIG. 4. σ^2/N as a function of $z \equiv 2^m/N$ for various values of N, on a log-log scale.

The two phase structure we have observed is due to competition between two different effects. First, there is an embedded periodic dynamics which results in strong positive correlations in the responses of the agents to subsequent occurrences of a given string of length m in G. For small m, it can be shown [6] that the system's response to any given m history in G is largely independent of its response to other m strings. In this phase, odd occurrences of a given m string in G result in a subsequent minority group whose size is close to $N/2$, while even occurrences result in very small minority groups. Moreover, the minority group that follows an even occurrence of a given m string in G is opposite that of the preceding odd occurrence of that same string. This gives rise to a "bursty" structure in L_1 with larger, order N excursions from the mean separated by smaller excursions of order $N^{1/2}$.

It is the response of the system to the even occurrence of strings that gives rise to the large deviations from the mean in L_1, and are responsible for the fact that, in the strategy-efficient, small m phase, σ (as well as $\Delta\sigma$) is proportional to N for fixed m. This dynamic also explains the flat conditional probability distributions such as those shown in Fig. 2(a). Consecutive (odd-even) occurrences of a given string produce opposite responses in the sequence of minority groups. Consequently, the conditional probabilities will be very close to 0.5 for all m strings, for a game with small enough m. Using a simple random walk argument [6], one can show that for values of $m \geq m_c$ this period-two dynamics ceases to dominate.

The second effect, and possibly the most remarkable feature of this system, is the emergent coordination among the agents' responses to different strings of length m which works, for large enough m, to reduce σ below the value it would have in a RCG. For m near m_c, the contribution of the periodic dynamics diminishes, and we uncover this remarkable emergent property which gives rise to an improvement in overall utilization of the resource. The region of greatest effective coordination (smallest σ) is when $z = 2^m/N$ is of order one. Coordination diminishes and σ approaches the RCG result as m increases beyond m_c. This can be qualitatively understood by recalling that each chosen strategy carries with it fixed responses to 2^m different strings of length m. Thus, the ranking of strategies by each agent must coordinate the agents' responses to 2^m different strings. As m increases, for fixed N, it becomes increasingly difficult for the agents to coordinate all of their responses (systemwide there are only N choices that can be made to try to satisfy 2^m conditions), and the system's behavior will look increasingly random as m increases for fixed N.

Despite this apparent random behavior for large m, the system is not in a Nash equilibrium for large m, nor generally for any other value of m. That σ is consistent with RCG at large m is not due to a Nash equilibrium mixed strategy that one would have in the RCG. In fact, there is increasing information in G as m increases [6]. There are many Nash equilibria in a minority game, but they are not achieved by the dynamics of this system.

It is also quite remarkable that σ^2/N lies on a universal curve as a function of the scaling variable $z = 2^m/N$ (Fig. 4). Since m_c is proportional to $\ln N$, we are led to the intriguing idea that for maximum coordination N should be roughly the same size as the *dimension* of the strategy space.

It is clear that the behavior of the system is qualitatively different for m above and below m_c. However, it is unclear whether that difference is the result of a singularity (and, thus, a bona fide phase transition at m_c) or a crossover effect, even in the limit $N \rightarrow \infty$ with z fixed. In Ref. [6] we show that information theoretic measures, including the entropy, appear to be nonanalytic at m_c (at least for large s), suggesting that a phase transition does exist, at least in the large s limit. It may also be that the phase change is accompanied by a change from a period to a nonperiodic (possibly chaotic) state [6,10].

In an effort to begin to understand the universality of our results, we have studied some models which are variations of the one described here, differing, for example, in the length of the history over which the strategies are evaluated, the nature of the publicly available information, or in some of the details of the way in which the agents choose among their strategies. While some details of the results change [6,11], the general structure remains the same. In particular, Fig. 4 is largely unchanged.

The model's general two-phase structure, with maximum utilization of resources at the phase transition (when the dimension of the strategy space is of the order of the number of agents playing the game), may well be a characteristic that transcends the class of simple models we have studies. Thus, the size of the available strategy space may be of practical significance for the structure of many systems such as financial markets and ecological systems.

Although the behavior we have elucidated is very intriguing, one must remember that there are many effects that may play a major role in specific systems and which could alter the emergent structure, fundamentally. For example, while this model is adaptive, it is not evolutionary. There is no discovery of new strategies by the agents, no mutation, no recombination, and no sex. Strong evolutionary dynamics may drastically alter the phase structure of the system. Nevertheless, any analysis of more complicated specific systems, which share the general competitive dynamics we have discussed here, must take account of the type of structure we have described.

Finally, and perhaps most importantly, our work raises the question: What really are the fundamental terms in which we ought to think about N-agent adaptive systems? For example, the fact that, for $m < m_c$, σ is so strongly dependent on the initial distribution of strategies suggests that a meaningful specification of σ must include a specification of the spread in the variance. This suggests the need for care in interpreting data from adaptive systems, or from their simulations. For the *cognescenti,* this suggests an intriguing analogy with spin glasses [12] in which the agents' random, but fixed, strategies are analogous to the frozen-in impurities found in a spin glass. In any case, it is certainly true that the phenomenon of frustration, which is at the core of glassy behavior, plays a significant role in competitive games of the sort described here.

The fact is that there is no well-developed epistemology for complex adaptive systems, and we are still quite unsure of what the important issues are or the most robust ways of characterizing the dynamics of such systems. But the study of simple statistical models, and the elucidation of the variety of behaviors which they manifest, can lead us toward a deeper understanding of how to properly frame the questions that we can sensibly ask, and sensibly answer, for complex systems.

[1] W. Brian Arthur, Am. Econ. Assoc. Papers Proc. **84**, 406–411 (1994).

[2] *Ecology and Evolution of Communities*, edited by M. L. Cody and J. M. Diamond (Harvard University, Cambridge, MA, 1975).

[3] *Coordination of the Internet*, edited by B. Kahin and J. Keller (MIT, Cambridge, MA, 1997).

[4] D. Challet and Y.-C. Zhang, Physica (Amsterdam) **246A**, 407 (1997).

[5] The dynamics of adaptivity are crucial to our results. It is thus essential that $s > 1$ so that the agents have more than one strategy with which to play. For $s = 1$, the game devolves into a game with a trivial periodic structure.

[6] R. Manuca, Y. Li, R. Riolo, and R. Savit, University of Michigan Program for the Study of Complex Systems Report No. PSCS-98-11-001 (to be published), available at http://www.pscs.umich.edu/RESEARCH/pscs-tr.html..

[7] The dependence of the results of the game on s are interesting and are discussed in Ref. [6]. However, the qualitative picture we present here obtains for $s \ll 2^{2^m}$.

[8] Classic references on market efficiency are written by E. F. Fama [J. Finance **25**, 383 (1970); **46**, 1575 (1991).] The sense in which the market is efficient here is subtle. The flat probability distribution means that the market is informationally efficient with respect to the strategies. But the market is not necessarily efficient with respect to the agents. Since the agents can switch their strategies, they could, in principle, use different strategies at different times and have a better than random success rate. In fact, in this phase the opposite happens and the agents' choices are maladaptive. See also Ref. [6].

[9] For different values of s there are systematic changes in the shape of this scaling curve, although the qualitative structure is similar. These are discussed further in Ref. [6].

[10] This may fall in the class of voting games discussed in D. Meyer and T. Brown, Phys. Rev. Lett. **81**, 1718 (1998).

[11] Y. Li, R. Riolo, and R. Savit (to be published).

[12] See, for example, *Spin Glasses*, edited by K. Fischer and J. Hertz (Cambridge University, Cambridge, England, 1991). See also Ref. [6].