8

Man and Superman: Human Limitations, Innovation, and Emergence in Resource Competition

Robert Savit,* Katia Koelle,† Wendy Treynor,‡ and Richard Gonzalez‡

1 Introduction

This conference is devoted to the theme of the design, prediction, and control of collectives. Many of the collectives that we are implicitly concerned with (and for good reason) are collectives composed of software agents or combinations of software and hardware agents. Collectives of agents that remotely gather information from distant planets and then transmit that information to Earth are one example. Of interest to the military are collectives of small, cheap sensors distributed on a battlefield or in a city that measure some aspect of local conditions and then relay that information to a central repository near a command center. Another example is a collection of sensors and actuators that control the flow of oil or electricity through a complex network by sensing local conditions and responding to them. One common architecture for the interaction of these local agents is through some sort of analogy with economic systems. Here it is supposed that the local agents compete for some scarce resource (bandwidth in the case of agents whose job it is to transmit information, or fluid pressure in the case of those agents whose job it is to regulate oil flow), possibly by a bidding mechanism or by some other strategic architecture that rewards agents for "buying low" or "selling high."

In all these cases, the local (software and hardware) agents rely on adaptive algorithms associated with that agent to make choices and execute actions. But in many cases the situation is more complicated. There will often be situations in which a human (or humans) can intervene in the workings of an otherwise automated, engineered system and change its behavior radically. Thus, a controller in Houston may decide to alter the mission or location of a set of remote sensors in space. A board of directors of an oil company may decide to intervene or countermand the actions of a network of local controllers on pipelines for the purpose of maximizing short-term

* Physics Department and Center for Theoretical Physics, University of Michigan, Ann Arbor, MI 48109

† Biology Department, University of Michigan, Ann Arbor, MI 48109

‡ Psychology Department, University of Michigan, Ann Arbor, MI 48109

profit. In still other systems of great importance the system may be entirely composed of human agents competing for a scarce resource. The archetypical example of such a system is financial markets, but there are others. Academics may compete for the scarce resource of tenure (or once tenure is secured, fame and a higher salary) by making a seminally important discovery before anyone else. Pharmaceutical companies (composed of humans) devote many resources to drug discovery in the hopes of patenting a lucrative (and, of course, beneficial) drug before their competitors.

These scenarios and many others like them naturally raise the question of how human agents behave when placed in competition for scarce resources. There are, generically, two ways in which such competition can proceed: conformity (or majority seeking) and innovation (or minority seeking). Examples of the dynamics of conformity include wars (the winner is usually the one with the bigger army, so it pays to be in the majority) and the acquisition of political power in a democracy. It is much better to belong to the party in control (i.e., the one that garnered the most votes and is thus the majority party). The examples cited in the previous paragraph, on the other hand, are largely examples of innovation or minority seeking. It only pays to discover and patent a drug if you're the first one to do it, and it's the one who *first* proves (and publishes) that great theorem who gets tenure.

Of course, in real systems it is the interplay between majority and minority seeking that gives them their richness and complexity. So, in war it is usually, but not always, the bigger army that wins. Sometimes David can beat Goliath by being innovative. In financial markets it pays to be innovative, but generally only if the crowd follows your example a little later. You want to buy IBM only if everyone else does—but does it the next day.

Nevertheless, in order to begin to develop a basis for understanding the complexities of competition in the real world, it is useful to look at very simple systems in which the naked dynamics of majority or minority seeking is exposed. In this paper we will concern ourselves with the dynamics of minority seeking, and, specifically, with the dynamics of minority seeking (or innovation) when it involves humans.

The minority game, first introduced by Challet and Zhang, and subsequently studied numerically and analytically[1] is a system in which pure innovation is rewarded. This model has also played a role in a number of other presentations at this conference. Detailed expositions of the model have appeared elsewhere, so I will not spend much time describing it. Most of the work done on the minority game has been in the context of either analytic treatments or computer simulations. In particular, there have, heretofore, been no controlled studies of how *humans* actually behave when faced with the need to place themselves in the minority. In this chapter I will discuss experiments we performed at the University of Michigan with human subjects playing the minority game.

[1] D. Challet and Y.-C. Zhang, *Physica A*, **246**, 407 (1997). R. Savit, R. Manuca, and R. Riolo, *Phys. Rev. Lett.* **82**, 2203 (1999). D. Challet, M. Marsili, and R. Zecchina, *Phys. Rev. Lett.* **84**, 1824 (2000). See also the references on the excellent Web site http://www.unifr.ch/econophysics/ minority.

2 Basic Features of the Minority Game

2.1 The Standard Minority Game

Definition

The simplest version of the minority game consists of N (odd) 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 in the minority group at that time step is awarded a point, and 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 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 the 2^m possible combinations of m 0s and 1s; 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 0s and 1s 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. An example of an $m = 3$ strategy is shown here:

Recent History	Predicted Next Minority Group
000	0
001	1
010	1
011	0
100	0
101	0
110	1
111	1

In the simplest version of the minority game, all strategies used by all 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. For its current play the agent chooses its strategy that would have had the best performance over the history of the game up to that time. Ties between strategies are decided by a coin toss. In the games we will discuss in this section, each agent has two strategies ($s = 2$). Because the agents 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, that is, in response to new entries in the time series of minority groups as the game proceeds. Note, however, that in these games agents must play with the strategies they were assigned at the beginning of the game. There is no evolution in these games. Evolution will be discussed later.

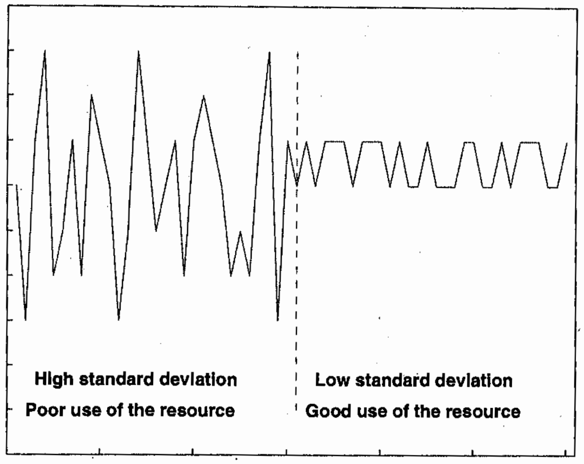

Figure 1.

Results—Analytic and Numerical

One of the interesting questions to address with the minority game is what the typical experience of the agents is. To put it another way, membership in the minority group is a limited resource. We might ask whether that resource is well used. If the minority groups are typically large (i.e., just less than $N/2$), then the limited resource is well used and a relatively large number of points are awarded to the agents in general. If, on the other hand, the minority groups are relatively small, then the limited resource is not well used, and not many points are awarded in toto. A convenient inverse measure of resource use is σ, the standard deviation of the number of agents in one of the groups, say, group 1. This may be thought of as a measure of inefficiency of the system at distributing resources. To see this, look at Figure 1. Because the game is symmetric, the expectation value of the number of agents in each group averaged over time is close to $N/2$.[2] If σ is large then the minority groups typically will be small, and if σ is small the minority groups will be large.

The basic result of the standard minority game is illustrated in Figure 2. Here we have plotted σ^2/N as a function of $z = 2^m/N$. Detailed discussion of this result can be found in the literature.[3] Here we only point out the most important features. First, all games lie on the same scaling curve. Second, if z is too small, then typically the system does poorly at allocating resources. If z is very large the system also does relatively poorly at allocating resources. In fact, for large z the aggregate behavior of the system, in the sense of resource distribution, is the same as if the agents were making their choices randomly. The dynamics here are somewhat different than sim-

[2] It is not quite N/2 in a given game because there are biases built into the game due to the particular (random) strategies that were assigned to the agents at the beginning of the game. This bias is not important for our discussion here.

[3] See note 1.

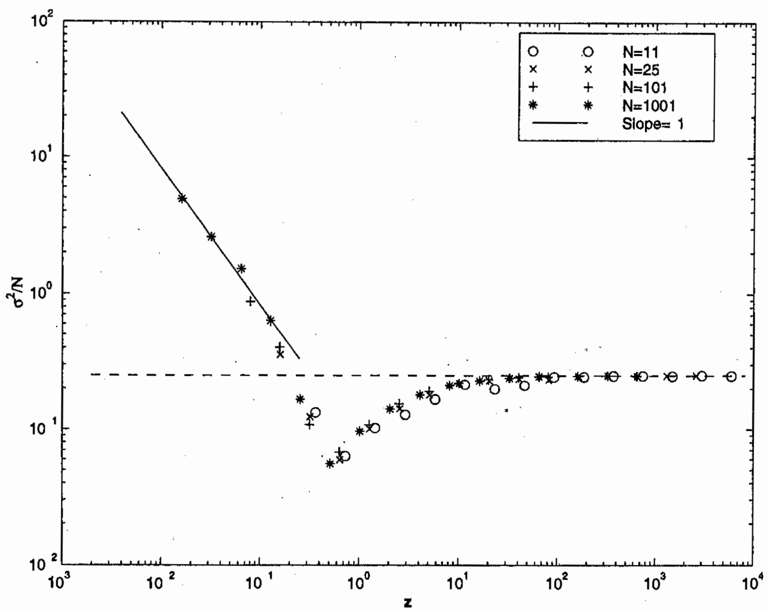

Figure 2.

ple independent random choices between the two groups, but the typical size of the minority group is the same as such a random choice game (RCG). For intermediate values of z we see emergent coordination in the agents' choices. That is, σ is significantly smaller than it would be if the agents were making uncoordinated or random choices, and the size of the typical minority group is relatively high. Note also that there is a minimum and an apparent cusp at a particular value of $z \equiv z_c$. At z_c there is a bona fide phase transition separating two phases with very different structure. Extensive discussions of this counterintuitive result can be found in a number of places in the literature.[4]

2.2 The Minority Game with Evolution

Before turning to a discussion of the human experiments, I want to describe one variation of the standard minority game that will be important for our subsequent discussion.[5] In the standard game, agents are assigned strategies at the beginning of the game and they retain those strategies during the entire game. However, in addition to adaptivity (being able to alter among established strategies in response to changes in the environment), agents in various systems can also manifest evolutionary capabilities. That is, agents can evolve their strategies, developing new ones

[4] See note 1.

[5] Y. Li, R. Riolo, and R. Savit, *Physica A* **276,** 234 (2000); *Physica A* **276,** 265 (2000).

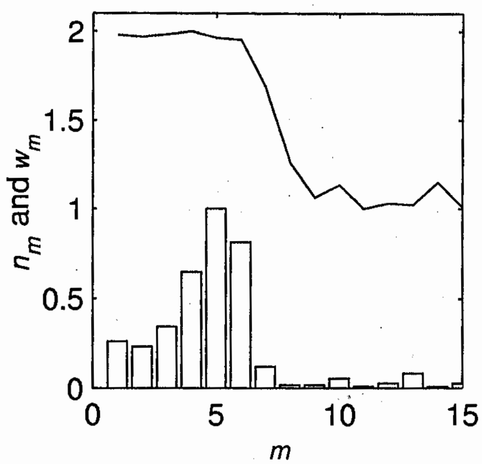

Figure 3.

and jettisoning old, poorly performing strategies. To model this dynamic, we modify the standard minority game in the following way: After playing the game for T time steps, a selection of the poorest performing agents are chosen, their strategies are discarded, and they are give two new (randomly generated) strategies. In a typical experiment, $T = 500$ time steps and half the agents from the bottom 20% of performers are chosen at random to have their strategies exchanged. The two new strategies can have a different memory, m, from the ones they are replacing so that the game now becomes one of competition among agents, each of whom may have a different memory. (Although one agent may play with strategies with different memory from another agent, the two strategies that a given agent has will have the same m as each other. That is, a given agent's two strategies always have the same memory during a given generation.) The results we describe here do not qualitatively depend on these parameter choices. This evolutionary process is repeated for a substantial number of generations (typically 20–50).

As one might expect, σ^2/N for the games with evolution is typically considerably smaller than σ^2/N for the nonevolutionary games. There is an additional important feature of these games illustrated in Figure 3. Here we plot some features of the population of agents after 50 generations in a typical minority game with evolution. In this case, $N = 101$. On the horizontal axis is m, the memory of the strategies of a given agent. Two types of data are plotted on this graph. The first, illustrated by the bar chart, shows the relative number of agents playing with a given memory. Note that almost all the agents play with memories of 6 or less.[6] The curve in this figure represents the average wealth of agents playing with a given memory. Note that

[6] The fact that there is a dramatic fall-off in the wealth per agent and in the number of agents playing with a given memory at $m = 6$ is not an accident. The value of m at the fall-off depends on N and is related to the value of z_c in the games played with fixed N and m. See Y. Li, R. Riolo, and R. Savit, *Physica A* **276,** 265 (2000), for more details.

the wealthiest agents (i.e., those that were most often in the minority group) were those playing with the *simplest* strategies, that is, the strategies that used the shortest memories. We will come back to these results later in this chapter.

3 The Human Experiments

3.1 The Setup

The human minority game experiments were conducted with volunteers from the University of Michigan and the Ann Arbor area. A room was set up containing 25 PCs. Partitions were placed between the computer screens so that players could not see what was on neighboring screens. Each player's screen contained two buttons that could be clicked with the mouse. Clicking the left button placed that player in group 1 and clicking on the right button placed that player in group 2. In addition, the screen contained a sequential list of which were the minority groups for all past time steps of the game. Each screen also contained a graph that showed that player's total winnings as a function of time. The history of minority groups and the graph of that player's winnings were updated after each time step of the game. No other information was available to the players, and no direct communication was allowed among players. Prior to the game being played, the rules of the game were explained to the volunteers. The volunteers were also told that they would be paid $10.00 for participating in the game and that players would receive $0.05 each time they were in a minority group. Finally, the players were informed that the player with the greatest accumulated wealth in the game would have its name entered into a drawing with nine other highest-scoring players from each of nine other games for a $50.00 bonus. This bonus was instituted to ensure that players would continue to focus and actively participate in the game. The volunteers were also told that they would be asked to complete a short questionnaire at the end of the game.

Each game was played for 400 time steps and took about 45 minutes. Players were allowed five seconds to make their choices in a given time step. A two-second grace period followed the five-second decision interval. A counter appeared on the screen so that the players would know the time remaining to make their choices. If at the end of seven seconds a player did not make a choice, the computer entered a random choice for the player. In different sessions, games were played with different numbers of players, N, ranging from 5 to 23.

3.2 Results

Among the most intriguing results of the computer simulations (and the analytic analyses) of the minority game is the observation of emergent coordination among the agents' choices. As shown in Figure 2, in games played with a fixed agent memory, for a range of values of memory, minority groups are typically larger than they would be if the agents were making random choices leading to greater generated average wealth. In the evolutionary versions of the game, in which agents' strategies

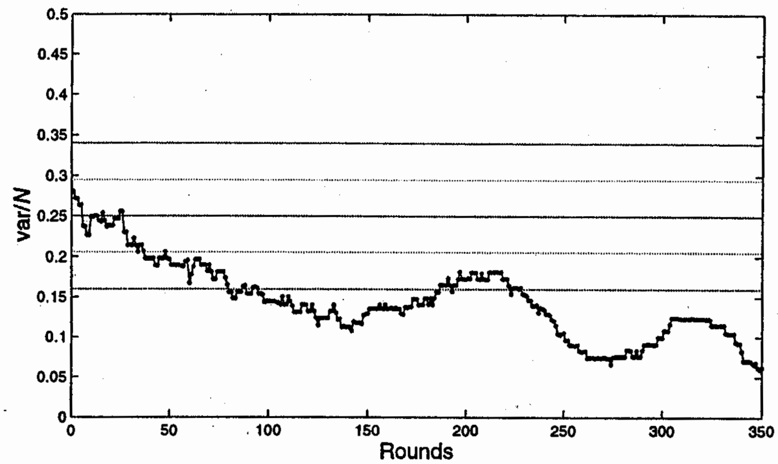

Figure 4.

and memories are allowed to evolve over time the emergent coordination is even more marked and the limited resource (minority group membership) is even better used. Or, to put it another way, more wealth is generated.

What happens when the game is played with humans? In Figure 4 we plot, for one of the games (in this case $N = 5$) σ^2/N as a function of time steps in the game. This result is qualitatively typical of the outcomes of our games. In this plot, σ^2/N is computed over a moving window of 50 time steps. The horizontal line at 0.25 indicates the expected value of σ^2/N for a random choice game (RCG). The lines at 0.2 and 0.3 (0.15 and 0.35) indicate one (two) standard deviation(s) about that expectation value under the RCG assumption. We see that at the beginning of the game σ^2/N has the value we would expect if the players were making their choices randomly and with equal probability. As time goes by, however, players' choices become more coordinated and the average size of the minority group increases, resulting in increased generated wealth. This is quite a remarkable finding. One might have supposed a priori that players, lacking information about the choices and strategies of other players, would have effectively made random choices of group membership. It appears, however, that this is not the case and that players effectively coordinate their choices based solely on the aggregate information provided by the time series of minority groups. One additional noteworthy feature is the occurrence of oscillations in σ^2/N with a period of 50 to 100 time steps. As we shall report in detail elsewhere,[7], these oscillations are the result of nonstationarity in the strategies of the players. In particular, players evolve their heuristics over time until a quasiequilibrium state is reached. Then one player, feeling that he can do better, will dramatically change his strategy. Because the game is endogenous, the other players must accommodate their strategies to the qualitatively new environment (i.e., the one generated by a qualitatively different strategy being played by one of the participants). During the period of

[7] R. Savit, K. Koelle, Y. Li, and R. Gonzalez, to appear.

adjustment overall efficiency of the systems degrades until a new quasiequilibrium state is reached, leading to oscillations in σ^2/N.

The results of Figure 4 are typical of results for most games played with various numbers of players. It is interesting to compare these results with similar ones generated by computer simulations. In Figure 5 we show some typical analogous results for computer simulation games played with evolution and with $N = 5$, 17, and 101 agents. In these examples the poorest agents are allowed to change their strategies every 20 time steps. Specifically, half of the least successful 30% of agents are chosen at random for strategy replacement. Their existing strategies are discarded and they are given two new strategies. The m value of the new strategies may be different than that of the old strategies, but the two new strategies share the same m value. Note the qualitative similarity to Figure 4, both in the decrease of σ^2/N over time and in the existence of oscillations.

The existence of emergent coordination in the human experiments, qualitatively similar to that observed in the computer simulations, is remarkable. Remember that the ways in which the computer agents choose their strategies is quite different from the ways in which humans develop and evolve heuristics. Nevertheless, the systemwide performance of the collectives is remarkably similar.

An equally remarkable observation concerns the behavior of individual agents and their strategies. Recall from Figure 3 that in computer simulations agents tended to adopt simpler strategies (as measured by the memory of the strategies) as the game progressed. Moreover, those agents with relatively simple strategies tended to fare better than agents with more complex strategies. Remarkably, with regard to simplicity, we find the same qualitative behavior in the human experiments.

Before presenting the results on individual players in the human experiments, it is useful to spend a moment discussing the definition of simplicity. There are at least two aspects of agent strategies that contribute to "simplicity." First is the amount of information agents use to make their decisions. In the context of the computer simulations, m is a simple measure of this component. The second aspect is the degree to which an agent's choice, given a certain amount of information, is deterministic. One simple way to measure this is to consider quantities like $A_m \equiv [P(1 \mid s_m) - 1/2]^2$. Here s_m is a specific string (of minority groups) of length m, and $P(1 \mid s_m)$ is the conditional probability that the agent in question chooses group 1 following the occurrence of the string s_m. If, given s_m, an agent's choice is completely deterministic then $A_m = 1/4$. If, on the other hand an agent's choice given s_m is completely stochastic, then $A_m = 0$. In the case of the computer simulations, agents with low memories (for example, in the case of Figure 3, agents with memories less than or equal to 6) generally have one of their strategies consistently more highly ranked than the other, so that $A_m \approx 1/4$ for all s_m associated with the memory of their strategies. On the other hand, agents with strategy memories above 6 have more closely ranked pairs of strategies. These agents bounce relatively often between these closely ranked strategies and so the associated A_ms for these agents tend to be relatively small.

In the human experiments we capture all the key strokes of the players, so we are able to impute effective strategies to the players and determine something about their

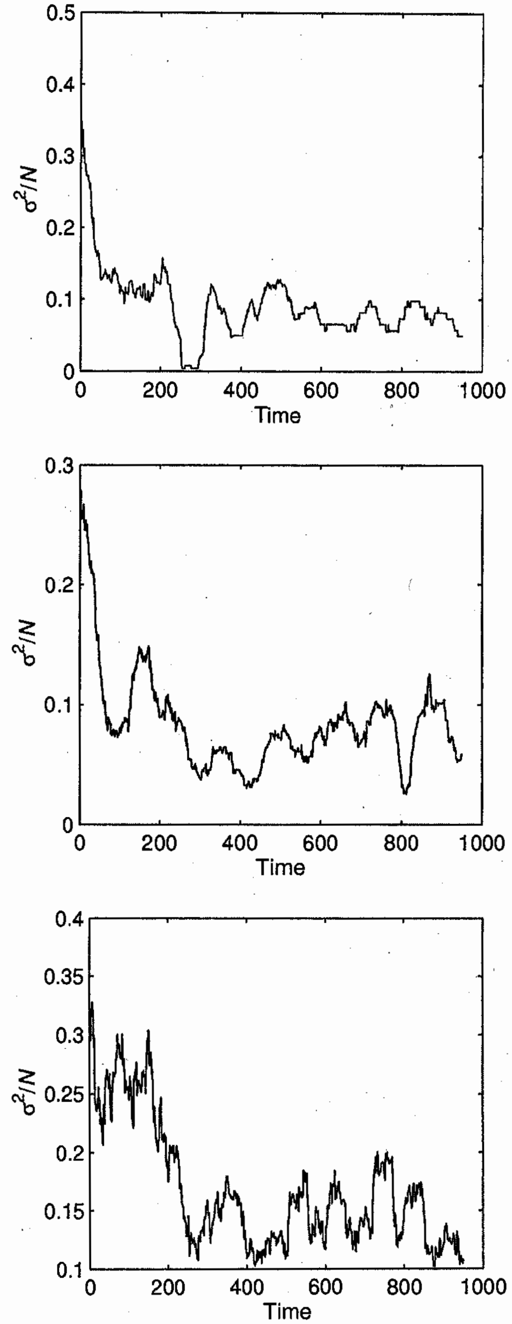
0.5
0.4
0.3
0.2
0.1
0
σ^2/N
0
200
400
600
800
1000
Time
0.3
0.2
0.1
0
σ^2/N
0
200
400
600
800
1000
Time
0.4
0.35
0.3
0.25
0.2
0.15
0.1
σ^2/N
0
200
400
600
800
1000
Time

Figure 5.

simplicities. In a forthcoming publication we will discuss this question in detail,[8] but here we present one simplified measure that indicates the relationship between simplicity and accumulated wealth. Consider

$$I_m \equiv \sum_{s_m} P(s_m) A_m$$

where $P(s_m)$ is the probability that the string s_m appears in the history of minority groups. If we assume that a player bases his choice of which group to join on information contained in no more than m lags of the time series of minority groups, then, I_m is a measure of the simplicity of the strategy. In particular, the larger I_m is the simpler the strategy is, in the sense that the strategy is more deterministic. Assume that the players in human experiments consider, at most, information from the last three time steps in deciding which group to join.[9] (We have also studied the m dependence of a player's strategy and have found that small m is also associated with greater wealth. This will be discussed in detail in a forthcoming publication.)

In Figure 6 we present a scatter plot of agent wealth versus I_m for a game with $N = 5$ players. In this figure, each shade of gray represents a different player. Each dot represents wealth versus I_m for that player over a 50-time-step window. Note first the obvious positive slope of the points indicating that, in general, there is a positive correlation between agent wealth and the simplicity (more properly, here, the degree of determinism) of an agent's strategy. Note also that there is a spread in the points of a given shading. In particular, one finds that, generally, points of a given shading that are further to the right (higher values of I_m) and consequently, in general, associated with greater wealth, occur later in the game. This is not too surprising because, generally later epochs of the game result in greater average wealth (see Figure 4). Finally, note that some players perform qualitatively better than others, and again, those players that perform better tend to adopt simpler strategies.

The general features of Figure 6 are found in nearly all the games we have studied for all values of N. In fact, the player who did the best out of all games was, I am sorry to say, an economist who pressed 1 all the time.

We have, therefore, two main observations. First, as in the computer simulations, humans playing the minority game demonstrate emergent group coordination. As the game is played, people learn about the actions of other agents through the medium of the publicly available aggregate information of the time series of past minority groups. Second, like the agents in the computer simulation of the minority game with evolution, there is a generally monotonic relation between the simplicity of a player's strategy and accumulated wealth.

There are several comments we would like to make about these observations. First, if we accept the observation that humans can generally pay attention to about

[8] See note 7.

[9] There is good reason to suppose that $m = 3$ is about the largest value of m that most people will use in making their decisions. There is a famous notion in psychology called 7 ± 2 that asserts that people typically can pay attention to between five and nine different pieces of information. Because $m = 3$ encompasses eight pieces of information, this is a reasonable value to take as a maximum number of lags that can be used to determine a person's choice.

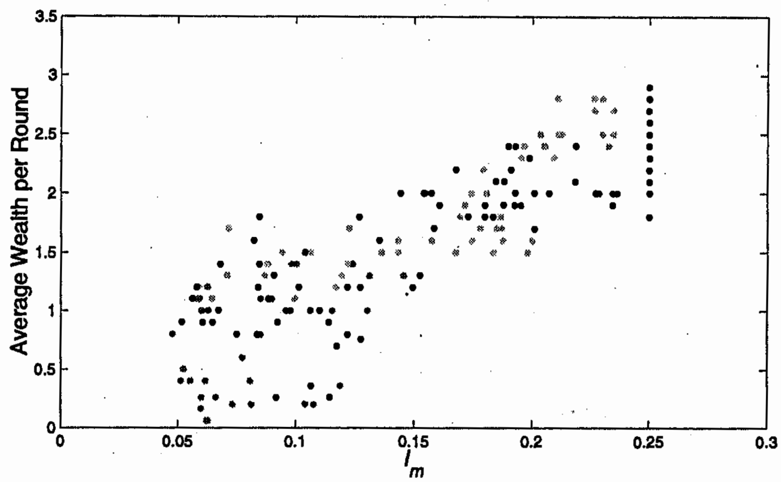

Figure 6.

seven different pieces of information (7 ± 2), then we expect that we can get good coordination of players' choices, and consequently good average wealth generation, if no player is forced to rely on strategies that use more than about seven pieces of information. Roughly, this means paying attention to no more than three previous time steps. If we further make the reasonable assumption that humans will pay the most attention to the most recent events, then we are led to deduce an upper limit on group size that can result in good wealth generation. Recall (see, for example, Figure 3) that in an evolutionary game in which different agents can use different memories, the wealthy agents have a memory, m less than some value m_t. If the processing power of the agents is not restricted, then $m_t = m_c - 1$. But in the case of humans, processing power is restricted, so we expect that $m_t \approx 3$. The effective dimension of the strategy space is therefore $\sum_{m=1}^{m_t} 2^m \approx 2^{m_t+1}$. For the best emergent coordination, the number of agents should be no greater than a number proportional to the effective dimension of the strategy space, with a proportionality constant of about 3.[10] Therefore, the number of agents that can be accommodated with a maximum memory of 3, in a game with agents with different memories, such that that collective achieves the best emergent coordination is $N \approx 3 \times 2^{m_t+1} \approx 40$. We thus expect that, for groups of fewer than about 40 players, we will see emergent coordination qualitatively similar to that which we have seen in our experiments to date. But for games played with groups larger than about 40 under the same circumstances, we expect that emergent coordination will degrade as group size increases.

Second, based on some of our exit polls, we noted that, although most people playing the game continued to be engaged during the entire duration of the game, some people became bored. When they became bored they tended to play very simple strategies, such as always joining the same group. This raises the interesting if

[10] See note 5.

somewhat speculative (note emphasis) idea that boredom may be an evolutionarily adaptive strategy. We may tend to get bored when we are engaged in an activity that is not particularly rewarding. Under at least some circumstances, namely those in which innovation can lead to success, the feeling of boredom and the attendant simple actions that it promotes may, in fact, lead to greater individual and group success.

Third, we have noted that there are oscillations in both the computer simulations and the human experiments. With regard to the human experiments, the period and nature of these oscillations tell us something about the nonstationarity of human strategic decisions. In particular, they go directly to the issue of exploration versus exploitation. When a player ceases being satisfied with his exploitation of the system he will begin to explore, and by that exploration (because the system is endogenous) he will upset the near equilibrium state of play, resulting in the performance oscillations we have observed. A closer analysis of these oscillations, their nature, and genesis, will shed light on the relationship between exploration versus exploitation in humans.

Finally, the research project of which these human experiments are a part involves direct confrontation of those experiments with computer simulations of similar games. We do not believe that all the subtleties of human behavior can be captured in simple computer games. However, there are clearly aspects of human and group behavior that have clear analogues in simple computer games. As we have seen in this chapter, taking the results of these computer games seriously affects the kinds of questions that we are led to ask about the human groups and individual behavior in those groups. This suggests a new epistemological thrust in psychology, and perhaps in the other social sciences. The *serious* confrontation with simple computational models can lead to a broader set of questions for the psychological community, and answers to those questions can generate a deeper understanding of individual and group behavior.

The primary observations reported here, namely the existence of emergent coordination and the correlation between strategy simplicity and agent wealth are particularly remarkable given the great difference in the ways in which agents choose their strategies in the human experiments compared to the computer simulations. The fact that both these features—emergent coordination and the dominance of simpler strategies—occur in such different settings, one computational and the other biological is striking. It suggests that there is a deeper, more general dynamic that underlies both these systems. This directly supports the implicit assumptions underlying this conference, namely that there are general underlying principles that are independent of the details of a particular collective and that can be used to guide the design and control of such collectives. More concretely, our observations point us in the direction of formulating at least tentative hypotheses concerning the general emergent group and individual behavior of agents, whether silicon- or carbon-based, competing for limited resources by being in the minority.

Acknowledgments. We are grateful to Yi Li for help with some of the calculations and data analysis. This work was supported, in part, by the University of Michigan's Office of the Vice-President for Research.