Adaptive Strategies in Game Theory

Alexander V. Outkin

December 3, 1998

Abstract

This dissertation is an attempt of expanding the domain of game theory into the sphere of evolving, potentially non-equilibrium systems. We especially focus our attention on studying the effects of local interactions, using automata networks as a modelling tool.

1 Introduction

Modern evolutionary game theory often treats evolution as a mere step by step repetition of the same, essentially static dynamics. At most, the agents would consider the past in their decision making. They rarely make any predictions of the future (except assuming it totally transparent). While this lack of desire to think about the future may be a perfectly reasonable assumption in the models of biological evolution, it may not be such in the world of humans. Moreover, a model considering the tactical value of decisions without thinking about the future falls somewhat short of the ambition of Game Theory to be The Theory of Strategic Interaction.

The necessity to think about the past and the future is usually eliminated by the equilibrium assumption. Indeed, in an equilibrium no decision making is necessary[1]. At every period of time a player needs to do exactly what he did just a moment before. For many equilibrium concepts, like for example Nash equilibrium, time is not an issue at all - the game is a one-time affair[2].

This is not to say that the equilibrium assumption is not useful. To the contrary, in many cases it does allow to produce important and meaningful results. Nevertheless, often the assumption completely fails. Human political, economic and business history can provide many illuminating examples.

The goal of this work is to begin expanding the domain and the methods of the Game Theory to the investigation of evolving, non-equilibrium systems.

Two questions arise immediately: first what features of the models (and hopefully of the world) create non-linearities and non-equilibrium properties;

[1] It is deemed necessary, however, in order to "arrive" at the equilibrium.

[2] It seems somewhat ironic, that evolutionary game theory is mainly preoccupied with the equilibrium concept.

second and subsequently how to adequately model the decision-making processes in this kind of systems and what is the "right" strategy for an agent to choose.

Let us start with question one. Many factors can contribute to the system being away from its equilibrium. One of them can be of course the inherently evolving nature of the system. When the agents, their attributes, and parameters of the system are in the constant state of flux, the most we can hope for in many cases is some kind of a temporary equilibrium.

However, even without modelling evolution explicitly, we may end up with a system that demonstrates non-equilibrium and sometimes chaotic behavior. I believe local interactions and learning by the agents are among the most important and relevant reasons for such behavior in the economic context.

The Chapters 2 and 3 concentrate on applications of the local nature of interactions and rely on automata networks as an investigating and modelling tool for game theory. Chapter 2 ("Cooperation and Local Interactions in the Prisoner's Dilemma Game," Outkin (1997)) devoted to cooperation and to a smaller extent to the endogenous formation of links between the agents. Chapter 3 ("Best Response Dynamics and Neural Networks," Haller and Outkin (1998)) is investigating the deterministic and stochastic best response play when interactions are local.

The second question appears to be even harder to answer. While it is obvious from the Chapters 2 and 3 that some strategies are better than the others in away-from-equilibrium situations, the general answer as to what makes a good strategy remains elusive. However, it represents a fascinating research agenda. Some attempts at it will be discussed in the current chapter.

The rest of the current Chapter 1 is organized as following: a general literature review, discussion of local interactions in population games, discussion of adaptive and learning strategies.

2 Literature

The literature selection for this chapter of the dissertation was a very difficult process. On one side there is a very limited amount of literature in economics on topics related to non-equilibrium systems, local interactions and adaptive strategies. These areas started to develop only recently in the economics framework, therefore the list is growing. On the other hand there is a very extensive literature on learning and automata networks (which can serve as an excellent model of local interactions) in computer science. Because of the aforementioned factors this literature review will not even attempt to cover all the relevant papers, but is intended rather to give a brief introduction into the topic. More detailed literature references, especially on local interactions, automata networks and cooperation, can be found in the Chapters 2 and 3.

2.1 Local Interactions, Groups and Automata Networks

2.1.1 Groups

Much of the evolutionary game theory literature on this subject draws on the framework developed by Kandori, Mailath and Rob (1991), which considers neither local nor group interactions but rather serves as a modelling prototype and a benchmark.

Canals and Vega-Redondo (1992) consider an “evolutionary process which proceeds in ‘parallel’ at different ‘hierarchical’ levels.” They analyze a generalization of KMR’s model. The population is partitioned into disjoint groups and evolution takes place at two levels. At a lower level the agents participate in pairwise contests within each group. The two basic components are imitation and experimentation. At each stage an agent has, with positive probability, an opportunity to change his strategy. The assumption is that an agent imitates the behavior that was most successful in the previous round. The experimentation dynamics induce an agent to mutate with positive probability away from the strategy selected at the imitation stage. At an upper level the groups themselves face some positive probability of adjusting the group’s strategy profile, a process essentially analogous to the imitation on the intra-group level. When the set of all groups subject to the adjustment is identified, the average payoff of agents in a group is compared to the highest average payoff among all groups. If the former is strictly lower, then the group strategy profile changes to match that of a group having the highest average payoff. The motivation for the model is that in the real world interaction does not takes place within the population as a whole. Economic agents communicate within different groups (families, neighborhoods, companies, . . .), and the influence of the outside population is largely channeled through intergroup dynamics. The population structure in the paper is not close to that of a human society, but the authors believe that it provides a good first order approximation.

It is shown that “Pareto dominant equilibria are the only long run states of the process if the number of groups is relatively large.” The authors also point out that “...the segmented kind of interaction contemplated here speeds up the desired convergence quite significantly” (compared to the KMR paper). An intuition for the importance of the large number of groups is provided with the comments for the next article.

Vega-Redondo (1995) develops a “hierarchic evolutionary model leading to long-run cooperation in the Prisoners Dilemma.” The model is similar to the one in Canals and Vega-Redondo (1992). The difference is that in the 1995 paper the underlying game is Prisoners Dilemma and the probability of a group dispersal/strategy profile adjustment may differ across groups. In the 1995 article the groups are subject to dispersal with positive probability, but the only difference is in interpretation, rather than in the model itself. The author assumes that the probability of dispersal is strictly positive for each group, for any strategy profile. After a group disperses, the empty niche is recolonized by individuals from groups having the highest average payoffs. It may be misleading to call this process a dispersal, because a dispersed group can in fact be imitated if it had the highest average payoff. The authors represent the dynamics as a Markov process. The main result is that if the number of groups is sufficiently large (for a given size of a group), then the whole population will play the cooperative strategy most of the time. The mechanism that generates convergence to the cooperative outcome here is the same as the one which ensures Pareto efficiency in the previous article. The probability that the whole population will move from a "bad" equilibrium to a "good" one decreases in the group size, while the probability of the opposite transition decreases with the number of groups. Specifically, when the whole population is in the "good" equilibrium and the number of groups is large relative to the number of individuals in each group, then a single group that switches to the "bad" equilibrium is likely to imitate a more successful cooperative group and switch back to cooperation before individual mutations in every group cause the whole population to switch to non-cooperation. Hence as the mutation rate goes to zero the competition among the groups becomes infinitely fierce. Authors consider interactions in the population as having local character, because agents interact only with other agents in their group. It seems that the large (relative to the size of a group) number of globally interacting groups undermines the local character of interaction.

Boyd and Richerson (1990) consider the effects of group selection on the equilibrium outcome. They show that "when there are multiple evolutionary stable strategies, selection among groups can cause the spread of the strategy that has the lowest extinction rate or highest possibility of contributing to the colonization of empty habitats."

The authors consider a metapopulation divided into subpopulations, or groups. Two strategies are available for agents. One is beneficial for the group as a whole, and the other is the deleterious one. At each step this strategy is adjusted as a result of social interactions inside of the groups and then as a result of migration, where a fraction of each group is replaced by individuals drawn randomly from the metapopulation. After migration occurs, the group will become extinct with positive probability. This will happen either through death of the members or through their dispersal throughout the metapopulation. The vacant habitats are assumed to be recolonized by the surviving population. The authors assume that the probability of extinction is a non-increasing function of the group's fitness.

To analyze this model analytically, the authors make two simplifying assumptions: "(1) extinctions are sufficiently rare, and subpopulations are large, and (2) the number of subpopulations is very large." These assumptions are made solely to simplify the analysis, since the authors claim that intuitively, larger subpopulations and weaker selections decrease the impact of group selection. The authors derive conditions under which the socially beneficial strategy will spread in the metapopulation and assess the probability that this will happen. Their numerical simulations are consistent with the analytical results even when extinctions are not rare.

Whittle (1982) considers a system of units, which may interact by forming a directed bond between any two of them. The state of the system at each moment in time is represented by random graphs, which need not be trees. Units can mutate between different states. This model is semi-spatial, and units are distributed with density r. "The process is . . . interesting in that it can show collective and crucial effects. If parameters [r in this model] are changed so that association is favored relative to dissociation, then at a critical point, very large connected clusters begin to appear; the phenomenon of aggregation." The author applies the general model to a two-role case, where "the roles are identified with those of trader and farmer." The economy then may be agrarian or mercantile (two polar cases). Also, the level of aggregation may vary. It is shown that the farmer/trader ratio changes discontinuously at a certain value of r, and this discontinuity is accompanied by aggregation.

The common feature of the above articles is that there is an individually inferior strategy that is beneficial for the group as a whole. Making group selection a more important factor than individual selection can make the Pareto-efficient equilibrium the long-run outcome of the evolution process. I think the main weakness of this approach is that the exact mechanics of how the group selection will finally prevail remains unclear. It also might be of interest to look on a model of group selection with the interaction inside of each group being local.

2.1.2 Local interactions

Ellison (1993) studies the rate of convergence. His model is essentially the same as in KMR, except for the local character of interaction. Players choose a best response to the actions of neighbors. He shows that local interaction generates dynamics that converge faster than the those in KMR.

Cooperation remains one of the main puzzles for Game Theory and some branches of biology. A new approach to this problem was suggested by Nowak and May (1992). They use a model similar to Ellison's. Interactions are local, but after each round a player is replaced by the most successful of his neighbors (including himself). This is a form of imitation dynamics. In computer simulations the authors found that the local interactions model with agents imitating the best paying strategy in their neighborhood produces interesting global behavior. Full non-cooperation does not seem to be a globally attractive point. Cycles in which cooperation coexists with non-cooperation have appeared in simulations. The underlying model is nonlinear and does not seem to allow representation as a threshold automata network.

Berninghaus and Schwalbe (1995) study the evolution of conventions in a finite population of boundedly rational players. They consider a model with agents interacting with a set of their neighbors, playing a coordination game. Players play pure strategies, and at the beginning of each period they choose the best response to the actions of their neighbors in the previous run. The authors show that this model can be analyzed in the framework of threshold automata networks. They confine their attention to analyses of evolutionary stability of conventions in one- and two-dimensional structures. In the one dimensional case they consider a model equivalent to a model with agents distributed over a circle. The authors show that the two possible strategy profiles with all agents playing the same strategy are the only fixed points of the dynamics given that all players have symmetric reference groups of the same size. It is also shown for the reference group of size two that the risk-dominant equilibrium is evolutionary stable of degree $n-1$ (n is the population size). For the two-dimensional case the authors discuss conditions that allow peaceful coexistence of conventions in case of von Neumann neighborhood . They relate the coexistence of different conventions to the dimensionality of the problem and to the degree of anonymity in the society. It is not clear whether the conclusion that anonymity makes possible coexistence of different conventions is justified, because they do not consider other possible two-dimensional types of neighborhoods.

Berninghaus and Schwalbe (1996) discuss applications of discrete iterations analysis to games with local matching. The model is the same as in the previous article, but a different mathematical apparatus is used - analysis of discrete iterations, which allows more general types of interactions to be considered, because it requires less information about the network. Discrete iterations theory can be especially useful when a network can not be represented as a threshold network. The two authors discuss the theory of discrete iterations (mostly convergence issues) and then apply it to games with local interactions and best response players. Concepts of global and local convergence are discussed. It is shown that in a heterogeneous population (where some players play best response and others imitate the most popular strategy in their neighborhood) dynamics will not converge globally. For a local Nash equilibrium, which is also a fixed point of the dynamics, the authors describe conditions sufficient to ensure local convergence to that point.

There is no significant body of literature on group interactions *per se* in the computer science (and related disciplines like engineering), but there is an absolutely vast number of articles, books, conference presentations on the theory and applications of various types of automata networks. Together with discreteness automata networks often imply local character of interactions. As Garzon (1991) put it: *Locality* is a fundamental restriction in nature. On the other hand, adaptive complex systems, life in particular, exhibit a sense of permanence and timelessness amidst relentless constant changes in surrounding environments that make the global properties of the physical world the most important problems in understanding their nature and structure." We can definitely apply this statement to the world of economics as well.

Coincidentally, Garzon (1991) serves as a excellent source of theoretical result in various types of automata networks, though potentially not including the most recent ones. Another excellent source on the applications of neural networks is Haykin (1998).

2.2 Adaptive Strategies and Learning

Perhaps, one of the first attempts to develop a decision-making framework for uncertain, non-equilibrium systems, can be found in Sun Tzu's "The Art of War". Although it is probably more than 2500 years old it is still considered to be of a great value by many military and business experts.

Page (1998) provides an interesting classification of various modelling scenarios and an attempt to unify some of the terminology, considering "...uncertainty, difficulty and complexity both as measures of environments and as analytical paradigms." The author also discusses various roles the institutions play in the society, especially in the confines of the complexity science paradigm.

A significant amount of research, considering the economy and the parts of it as complex adaptive systems (not limited to the automata networks) has been conducted in Santa Fe Institute. Part of this work is gathered together in two (so far) volumes of "The Economy as an Evolving Complex System" (I and II).

Sonsino (1997) discusses learning in the sense of ability to recognize patterns and convergence to the Nash equilibrium. He shows that under reasonable assumptions on rationality of the agents and on the structure of the underlying game, one can prove convergence to Nash equilibrium. The author, considers, however, only a game between two players; thus the results may not necessarily apply to a population game with local interactions.

Given its sheer size, the learning literature in computer science cannot be comprehensively reviewed here, given the size of this work. Again Haykin (1998) provides a first brief introduction to many. Other good sources include Muller (1996), Goles and Martinez (1990) and Suykens et. al. (1996).

Related literature includes articles and books on different types of machine learning (supervised and unsupervised learning, reinforcement learning), also various types of learning in (and by) networks (neural networks, Bayesian networks, etc.). A broad overview and literature references can be found again in Haykin (1998). On a more detailed level and for more specific discussion and applications, one can refer to Muller (1991), Maes (1989), Bay and Stanhope (1997).

3 Local Interactions

The reasons for non-equilibrium dynamic phenomena, probably cannot be fully explained in any monograph of a finite size. However, let us point out that the inherently local nature of interactions in the real world seems to attribute immensely to these phenomena, mainly for the following reasons:

Local interactions increase dramatically the state space of the system.

Local interactions are related to the creation of groups and coalitions. They not only enlarge the state space of the system but also may add another level of complexity to the agent's choices: whom to play with, what groups to be a member of, etc.

Local interactions has a built-in lack of a complete up-to-date information about the state of the world.

Local interactions can introduce non-linearities in the dynamics of a system, that would be linear when the interactions were global.

Further, when interactions are local, the system's dynamics often may no longer allow description by the standard probabilistic methods. As will be shown in Chapter 2 even a totally deterministic finite system of locally interacting agents can demonstrate dynamics that is seemingly chaotic and/or random.

It seems necessary to take this sort of uncertainty and incomplete information into account when we seek to build foundations for our thinking and for modelling economic agents. A proper way to do it may be very different from assigning somewhat arbitrary probabilistic weights to the future states of the world, because these weights generally will be affected by the (and especially so local) interactions between the agents.

4 Adaptive Strategies

As a result of this disequilibrium, incomplete information view of the world, comes a conclusion that a **really rational** decision-maker should not always rely on an apparently rational (or boundedly rational) strategy like best response. "What follows can not be grasped even by skilled calculators, much less by ordinary people."[3]

But what are the alternatives? For example, an agent can try to investigate all the possible scenarios of what can happen in the future (perhaps using historic information) and then choose the however defined best one. This can prove even less fruitful for two fundamental reasons:

1. To predict the future evolution of the system we need to have a model of other agents in the population. The true model of other agents' decision making may be impossible to develop and even if it does exist this model itself may be a subject of evolutionary change. This alone, even in a population of two agents makes accurate predictions practically impossible. There is however another reason:

[3] "The Essential Tao: An Initiation into the Heart of Taoism Through the Authentic *Tao Te Ching* and the Inner Teachings of *Chuang-tzu," 1991,* translated by Thomas Cleary.

2. Say we decided to optimize n periods ahead. Then this becomes an optimization problem with a huge state space that may be impossible to solve analytically or numerically, unless one just goes through all the possible alternatives. A related problem is that at every given step, the true value of what an agent has achieved already is impossible to evaluate exactly, because it is dependent on the unknown future. Best response in a Prisoners' Dilemma Game may serve as an excellent example of a strategy that may give the best possible payoff today but will surely fail tomorrow.

There may be a solution falling between these two extremes:

1. From the point of view of an agent, it still may be better to use some heuristic strategy that can provide a good payoff even though an agent may not know why. [4] While these strategies may not demonstrate an excessive understanding of the problem, the experience embedded into them may prove to be sufficient.

2. Another, far more exciting opportunity is that the population as a whole, may learn (with or without learning on the part of an individual agent). This then essentially becomes a collective problem solving task. A goal in this case will be to develop a set of relatively simple agents' strategies that will provide reasonably good individual payoffs and global performance of the system as a whole.

In the context of the second statement above it becomes important not only what the player does, but also how he interacts with the other players. Is there information sharing? What are the sets of players who share information, whom the player interacts with and how those agents are chosen, etc.?

While this aspect of game theory has gotten very extensive attention in previous work, still a lot remains to be investigated and understood. A very big part of what the player does is actually not just choosing the strategy to play, but choosing whom to play it with. Importance of this decision making dimension is recognized by the business-oriented literature, but the theoretical development has a long way to go.

There is no firmly established definition of what adaptation means, at least not in Game Theory. We will mean under adaptation and by an adaptive strategy a strategy that uses rational as well as heuristic rules (rules of thumb for example) and changes them according to the player's experiences in the game. It would be very interesting to consider a strategy that can combine short-term decisions with the long-term ones, which can be a subject of a future research.

There is not really a clear-cut distinction between rational and heuristic rules. For example if two players arrived at the same rule, one by making

[4] For example experienced surfers claim that every seventh wave in the ocean is a big one.

calculations and another just by using a wild guess, does this make the rule rational or irrational? Best response strategy may serve as a good example: being as it claims "best response" it is also equivalent to imitating a certain fraction of the player's neighbors.

5 Building Adaptive Strategies and Adaptive Agents

After we had concluded that bounded rationality (and perhaps unbounded too in certain circumstances) can not guarantee reasonable performance by the economic agents, the question arises: How can we improve it? How should we build adaptive strategies and adaptive agents?

There is fairly little research in economics on the subject. However, there is a huge amount of work on learning and adaptation in computer science and related disciplines.

Usually the decision making process of the economic agents is modelled using the utility concept. Utility theory also allows to achieve a great degree of understanding of human behavior, at least in an economic context and also to build analytically tractable models. However, in many instances it does not confirm to the common sense of even the scientists who subscribe to the idea. Yet, often it is important to have a benchmark or some way to evaluate performance of the agents, to allow them to learn using the feedback from their own and others' past experiences. Utility, profit, or some other perhaps similar measure, can serve this purpose well.

In addition or sometimes instead of using the utility concept, we can consider why people act. Generally speaking, we can define the driving motives of people as three broad categories:

- People want to achieve goals.
- People often do what is prescribed by the society, circumstances, friends, their own likes and dislikes, etc.
- People want to feel good about their actions or the results of those actions.

Clearly, the items above are not necessarily mutually exclusive, but a classification of this kind can be useful for modelling the human behavior. Therefore in order to adequately specify the strategies of adaptive agents, we have to make goals being part of the strategy specification. Another very important part is learning by the agents - in what circumstance choose which strategy, which goal or subgoal and whom (for the extent of agent's ability to choose) interact with.

The goal of this chapter is not so much to build any models or derive any new results, but rather to provide a general discussion and introduction to what will follow. This discussion, nevertheless, is easier to develop when it is based on some more or less formal framework, which the next section will attempt to develop.

6 An Example of a Formal Framework

There are many possible ways and tools available to formally describe and investigate this kind of issues. One is direct simulation, another one is paved with the apparatus of Markov Random Fields. The beauty of Automata Networks is that they allow a realistic enough modelling framework that affords the ease of interpretation and applications and even theoretical results in certain cases.

We will start by introducing a few definitions to establish a formal framework for our discussion using the concept of automata networks. The area of automata networks has, perhaps, not yet matured enough. Therefore, it is difficult at times to know what exactly a particular object means for different people. Nevertheless it seems that the standard definition of Automata Networks is somewhat limited for our purposes. We will start therefore with introducing a notion of Generalized Neural Networks.

A more detailed discussion can be found in Haller and Outkin (1998). We will first introduce a formal definition of a Generalized Automata Network (GAN), and then the standard definition of an Automata Network and a Neural Network (NN). Throughout, I is a finite or infinite set.

Definition 1 *A* **Generalized Automata Network** *on I is a triple*

$$A = (G, \Sigma, \Gamma, (f_i, i \in I)) \text{ where:}$$

$G = (I, V)$ is a (undirected) graph on I with connections given by the set $V \subset I \times I$. We assume G to be locally finite, which means that every neighborhood $V_i = \{j \in I : (j, i) \in V\}$ is finite[5]*.*

Σ is the set of states of a node of the graph G. It is usually assumed to be finite. A state of a node i is denoted by s_i.

Γ_i is the information set of a node i of the graph G. $\Gamma = (\Gamma_1, ...\Gamma_N)$. Note: I am not sure this is a sufficient amount of elaboration.

$f_i : \Sigma^{V_i} \bigotimes \Gamma_i \to \Sigma$ is the transition function associated with the vertex i. The evolution in discrete time t of the global state $s(t) \in \Sigma^I$ is governed by the global transition function $F : \Sigma^I \bigotimes \Gamma \to \Sigma^I \bigotimes \Gamma$ obtained as composition of all the local ones. Each f_i is also called an automaton.

Definition 2 *An* **Automata Network** *on I is a triple*

$$A = (G, \Sigma, (f_i, i \in I)) \text{ where:}$$

$G = (I, V)$ is a (undirected) graph on I with connections given by the set $V \subset I \times I$. We assume G to be locally finite, which means that every neighborhood $V_i = \{j \in I : (j, i) \in V\}$ is finite[6]*.*

[5] As can be seen from the definition, $i \notin V_i$.
[6] As can be seen from the definition, $i \notin V_i$.

Σ *is the set of states of a node of the graph G. It is usually assumed to be finite. A state of a node i is denoted by* s_i.

$f_i : \Sigma^{V_i} \to \Sigma$ *is the transition function associated with the vertex i. The evolution in discrete time t of the global state* $s(t) \in \Sigma^I$ *is governed by the global transition function* $F : \Sigma^I \to \Sigma^I$ *obtained as composition of all the local ones. Each* f_i *is also called a (memory-less) automaton.*

Definition 3 *A* **Neural Network** *or* **Threshold Automata Network** *is a particular type of automata network. Its state space is binary, e.g.* $\Sigma = \{-1, 1\}$. *Here we assume for convenience that* $\Sigma = \{0, 1\}$. *The network's transition function and, implicitly, its graph are based on a weight structure, given by a matrix* $A = (a_{ij})$. *Namely: every arc* $(i, j) \in V$ *is associated with a real number* $a_{ij} \in \mathbb{R}$ *which represents its weight. If i and j are not neighbors, we put* $a_{ij} = 0$. *The transition function takes on the following form:*

$$s_i(t+1) = L\left(\sum_{j=1}^{N} a_{ij} s_j(t) - b_i\right) \tag{1}$$

where b_i *is a threshold and the function* $L(\cdot)$ *is given by:*

$$L(x) = \begin{cases} 0, & x \leq 0 \\ 1, & x > 0 \end{cases} \tag{2}$$

An automaton described by (1) is called a **neuron** *or* **threshold automaton**.

Game-Theoretical Interpretation. We will consider Σ as the strategy set of a player, with an individual strategy of a player i denoted by s_i, $s_i \in \Sigma$.

We will denote the payoff to player i at time t by $\pi_i(s_i, s_{-i}(t))$, where $s_{-i}(t)$ is the distribution of strategies in V_i at time t. For example in 2x2 games it will be determined by the game matrix:

$$\begin{array}{ccc} & \sigma_1 & \sigma_2 \\ \sigma_1 & a, a & b, c \\ \sigma_2 & c, b & d, d \end{array} \tag{3}$$

In most of the following discussion of best response and imitation we will be using an example of the prisoners dilemma game, represented by the following matrix:

$$\begin{array}{ccc} & \sigma_1 & \sigma_2 \\ \sigma_1 & 0, 0 & b, 0 \\ \sigma_2 & 0, b & 1, 1 \end{array} \tag{4}$$

The transition function f_i and F correspondingly may (or may not) depend on the players' payoffs $\pi(s_i, s_{-i})$ For an example of applying Automata Networks in a 2x2 game see Haller and Outkin (1998).

7 General Meaning of Adaptation and Specific Adaptive Strategies

7.1 Decision Making by Adaptive Agents

Consider the decision problem of an individual agent. Consider the player's payoff $\pi(s_i, s_{-i})$. As one can easily see that payoff may be a very complicated function, especially in a more general situation than 2x2games. Using the terminology of landscapes (Kauffman, 1993, 1995), the landscape defined by the $\pi(s_i, s_{-i})$ may be very ragged. Therefore achieving the maximum of $\pi_i(s_i, s_{-i})$ may be difficult or impossible. To name just a few reasons: that maximum may not be easy to calculate and even if it is reached by an agent it may be unstable from the point of view of the global dynamics. Also stability against a single perturbation may not be enough. What we are really interested in are the regions in the state space that are stable against constantly arriving noise. Using terminology of Young (1993), we can try to divide the state space into the classes of stable states. We can consider as an implicit or explicit goal of the population of players to reach one of these regions, preferably one with a sufficiently high degree of appropriately defined mutual satisfaction. One way to quantify this criteria is by looking at the average payoff in the population.

A goal of the population of agents then would be to identify the individual rules that will allow the independently acting agents to achieve one of these acceptable outcomes. A population can potentially improve the speed and efficiency of this process by adaptation of new rules and by allowing the agents to communicate with each other in the process.

The above goal may not sound very easy to attain, it may also involve excessive computational requirements. We know, however, that the explicit understanding on the part of agents of what is going on is not absolutely necessary for successive collective problem solving. For example ant colonies perform very sophisticated tasks very efficiently [7] , yet one probably would not talk about intelligence of an ant. Human organization is another example. For instance, even if there is not much understanding of how exactly the markets work and, most of all, how they achieve at times an apparently quite efficient state, markets continue to exist despite the lack of understanding on the part of their participants. Furthermore, the market mechanism has actually become a tool for a computer scientist to solve certain optimization problems.

In order to get a foothold on the meaning of adaptation in the economics context we will compare best response and imitation in the local interactions context.

[7] "Ant algorithms" have been successfully employed to solve complex optimization problems in places ranging from Los Alamos Labs to brokerage houses.

7.2 Best Response Versus Imitation.

The best response rule is defined in order to introduce the Nash equilibrium. This I believe is one of the main shortcomings of best response, especially in an evolutionary setting, for it is essentially an equilibrium strategy that need not perform well away from an equilibrium.

7.2.1 Reasons for Imitation

Imitation appears to be a much better strategy from an individual and from the population point of view. For example it can produce some degree of cooperation in the Prisoners' Dilemma game (Nowak and May (1993), Outkin (1997)). The question is what is so special about the imitation dynamics that it induces cooperation? What other strategies will provide similar adaptive capabilities for an individual and for the society as well?

Of course, immediately a question arises why to select imitation? Why is it reasonable to assume that people imitate? It seems to be stupid to do in certain situations. We can ask another question. Why does a fisherman wait in one place for hours for a fish to bait instead of running all over the lake? For a fisherman we can say that he knows from his own and others' experience that this is the strategy that works best. An answer for imitation is similar. We can interpret imitation as learning from others' and one's own experience and as a related process of exchange of information. Imitation is a strategy that perhaps can be acquired as a result of experience.

Another consideration that may give rise for imitation without explicitly using it is the fact that people do not always know the structure of the game or what would be an equilibrium strategy. Therefore players may need to learn important features of the game. This may include using their own and others' experience, thus producing some hidden form of imitation.

A further consideration in favor of imitation or other adaptive strategies is that in a game modelled as an automata network, the resulting dynamic behavior can be very complex. Yet, the use of best response relies heavily on the assumption that a player can possibly understand the dynamics of the whole population. In general this is impossible, as long as we assume that an agent and his neighborhood is only a **part** of the whole system. This statement becomes clear if we look at a model with local interactions. Even if the rules governing behavior of the players are simple, the dynamics of the system may be very complex. It should be clear that best response is difficult or impossible to use not merely because of the complexity of computations involved. It is far more than that! From computer science we know that in a general case of an automata network there is no shorter way to predict dynamic behavior of the network than the direct simulation of it. In other words: Even if we have a collection of relatively simple automata, the dynamics of the system as a whole can - and usually will - be far more complex than behavior of any of its components. Therefore a component of a system can not possibly (in general) predict the behavior of the system. Therefore, even in a simple deterministic model, it is practically impossible for an agent to take into account all the strategic implications of his choices. The task becomes even more difficult it we introduce uncertainty or incomplete information, inherent in the local interaction models.

Therefore it should not be very surprising that a strategy like imitation without much claim on rationality can often outperform in an individual sense - and on the level of the population as well - another apparently far more rational strategy.

7.2.2 Evolutionary Advantage of Imitation over Best Response.

Let us consider a population comprised of best response players mixed with best performance imitators, engaged in the Prisoners' Dilemma game. Clearly, the dynamics of the system will be driven by imitators, as the best response players will always choose defection. Therefore, to investigate the equilibria set of this population we need only take behavior of the imitators into account, and consider the rest of the nodes being always zero. Then we can reformulate the Theorem 1 in Outkin (1997), with the proof being essentially the same as in the mentioned paper.

Theorem 1 *Consider a population of best response and imitative players, engaged in the Prisoners' Dilemma Game on an arbitrary graph G, representing a neighborhood system. In a mixed equilibrium (if it exists) the highest payoff achieved by a best response player can not exceed the highest payoff achieved by an imitator.*

The proof is essentially the same as in Outkin (1997). The existence of the equilibrium will depend on the number and location of the best response players, however. A trivial example of an equilibrium of this kind would be a population on a lattice split into to halves - one populated by best response players and the other by imitators, all of whom cooperate, then if b is small enough the described allocation will be an equilibrium). It also follows from Outkin (1997) that the only type of equilibrium in a one-dimensional case will be a population consisting of clusters of cooperators and clusters of non-cooperators.

One can easily conjecture from the above that in an evolutionary setting imitation will have advantage over the best response, that may depend, however, on particular details of the dynamics and of evolutionary selection. Inferiority of the best response should not be surprising especially in the light of the conclusion in Haller and Outkin (1998) that the best response is a special case of a majority or a minority imitation, which intuitively is a less sophisticated type of behavior then performance imitation (imitation in text).

7.2.3 Interpretations of Imitation Dynamics

There are basically three possible interpretations of imitation dynamics: 1) imitation *per se*; 2) a replicator dynamics when a cell is replaced by the offspring of its more successful neighbors and 3) exchange of information. In the latter quality (exchange of information) imitation can be considered as a very sophisticated strategy. From a justification of the performance imitation as a strategy that had worked for a neighbor of a player in a similar set of circumstances and therefore may work for that particular agent as well, we conclude that performance imitation may allow adaptation even in situations when agents have very little information about the game.

7.2.4 Modifications of Imitation Dynamics

Imitation as we have discussed above is a very interesting strategy, but difficult to analyze. There are many types of dynamics in networks for which no prediction of the outcome of the dynamic process is possible, short of running a simulation. It is plausible that imitation belongs to this type of dynamics.

In certain situations imitation gives very good results for individual players and for the population as a whole. Our goal is to analyze what features of imitation provide these results. We can try to identify a few main features of imitation that can induce cooperation in the population. Further we will try to approximate some of these features by a simple strategy, that has more "rational" appeal than standard imitation and yet, will support cooperation in the Prisoners' Dilemma Game.

The main features of (modified?) imitation as I see it:

- The strategy of a player in the next period depends not only on strategies and payoffs in his neighborhood but on a larger set of players (the neighbors' neighbors)
- In a certain sense the players take the future into account - even though they base their choice of strategy only on what happened today. The interpretation of imitation in this context is that the players imitate their neighbors, because players themselves are not quite sure what strategy will work, but observing their neighbors whose circumstances are not so different, they assume that the best (or merely adequate) strategy in the neighborhood today will also work tomorrow. In other words the rationale for imitating the neighbors lies in the fact that the environment they face is similar to the imitator's.

The next section considers a model of sophisticated best response that takes these features into account.

8 Conclusions

We argue that in an uncertain and non-stationary world, a reasonable choice is to use a simple adaptive strategy, rather than a strategy that will attempt to model the world precisely. We claim that in many situations the imitation dynamics serve as a much better adaptation tool than a best response dynamics. We discuss two other strategies that seem to perform well in cases where best response is inadequate: sophisticated best response and imitation with incremental improvements. We further discuss an example of a model with explicit information sharing and coalition formation. We claim that in such a setting a group of companies (agents) can solve a problem in a far more effective way than any single one of them.

REFERENCES

Bay, J.S., and J.D. Stanhope (1997): "Distributed Optimization of Tactical Actions by Mobile Intelligent Agents," Journal of Robotic Systems, 14, no. 4, 313-323.

Berninghaus, S. K., and U. Schwalbe (1996a): "Evolution, Interaction, and Nash equilibria," Journal of Economic Behavior and Organization, 29, 57-85.

Berninghaus, S. K., and U. Schwalbe (1996b): "Conventions, Local Interaction, and Automata Networks," Journal of Evolutionary Economics, 6, 297-312.

Boyd, R., and P. J. Richerson (1990): "Group Selection among Alternative Evolutionary Stable Strategies," Journal of Theoretical Biology, 145, 331-42.

Canals, J., and F. Vega-Redondo (1992): "Multilevel Evolution in Games," working paper.

Ellison, E. (1993): "Learning, Local Interaction, and Coordination," Econometrica, 61, 1047-71.

Eshel, Illan, Emilia Sansone, and Avner Shaked (1996a): Evolutionary Dynamics of Populations with a Local Interaction Structure, mimeo.

Eshel, Illan, Larry Samuelson, and Avner Shaked (1996b): Altruists, Egoists and Hooligans in a Local Interaction Model, mimeo.

Garzon, Max (1995): *Models of Massive Parallelism.* Berlin: Springer.

Golez, Eric, and Servet Martinez (1990): *Neural and Automata Networks.* Dordrecht: Kluwer.

Kandory, M, G. Mailath and R. Rob (1991): "Learning, Mutation, and Long Run Equilibria in Games," Econometrica, 61, 29-56.

Kauffman, S.A. (1993): "The Structure of Rugged Fitness Landscapes," in The Origins of Order: Self-Organization and Selection in Evolution. New York: Oxford University Press.

Kauffman, S.A. (1995): *At Home in the Universe.* New York: Oxford University Press.

Maes, P. (1989): "How to do the Right Thing," Connection Science, 1, no. 3, 291-323.

May, R. M. (1994): "Spatial Chaos and its Role in Ecology and Evolution" in Frontiers in Mathematical Biology, ed. by Simon A. Levin. New York: Springer-Verlag.

Muller, Jorg P. (1991): *The Design of Intelligent Agents.* Berlin: Springer.

Nadal, Jean-Pierre (1994): "Formal Neural Networks: From Supervised to Unsupervised Learning" in Cellular Automata, Dynamical Systems and Neural Networks, ed. by Goles, Eric, and Servet Martinez. Dordrecht: Kluwer.

Nowak, Martin A., and Robert M. May (1993): "The Spatial Dilemmas of Evolution," International Journal of Bifurcation and Chaos, 3, 35-78.

Nowak, Martin A., Sebastian Bonhoeffer, and Robert M. May (1994): "More Spatial Games," International Journal of Bifurcation and Chaos, 4, 33-56.

Page, Scott E. (1998): "Uncertainty, Difficulty, and Complexity," SFI Working Paper.

Peretto, P, M. Gordon, and M. Rodriguez-Girones (1992): “A Brief Account of Statistical Theories of Learning and Generalization in Neural Networks” in Statistical Physics, Automata Networks and Dynamical Systems, ed. by Eric Goles and Servet Martinez. Dordecht: Kluwer.

Rojas R. (1996): Neural Networks: A systematic Introduction. Berlin: Springer-Verlag.

Suykens, Johan A.K., Joos P.L. Vandewalle, and Bart L.R. Moor (1996): *Artificial Neural Networks for Modelling and Control of Non-Linear Systems.* Dordecht: Kluwer.

"The Economy as an Evolving Complex System II " (1997): (ed. by W Brian Arthur, Steven N. Durlauf, and David A. Lane). Reading, Massachusetts: Addison-Wesley.

Vega-Redondo, F. (1993): “Competition and Culture in the Evolution of Economic Behavior: A Simple Example,” Games and Economic Behavior, 5, 618-31.

Vega-Redondo, F. (1995): “Long-Run Cooperation in One-Shot Prisoners Dilemma: A Hierarchical Approach,” mimeo.

Whittle, P. (1982): “Criticality and the Emergence of Structure,” in Evolution of Order and Chaos, ed. by H. Haken. New York: Springer-Verlag.

Wynne-Edwards (1962): Animal Dispersion in Relation to Social Behavior. Edinburgh: Oliver & Boyd.