

Genetic Algorithms and Economic Evolution

THOMAS RIECHMANN

Diskussionspapier Nr. 219

ISSN 0949-9962

Dezember 1998

JEL classification: C63, D73, D83

keywords: learning, computational economics, genetic algorithms, evolutionary dynamics

Thomas Riechmann, Universität Hannover, Fachbereich Wirtschaftswissenschaften, Institut für Volkswirtschaftslehre, Abteilung Wachstum und Verteilung, Königsworther Platz 1, 30167 Hannover, e-mail: riechmann@vwl.uni-hannover.de

This paper tries to connect the theory of genetic algorithm (GA) learning to evolutionary game theory. It is shown that economic learning via genetic algorithms can be described as a specific form of evolutionary game. It will be pointed out that GA learning results in a series of near Nash equilibria which during the learning process build up to finally reach a neighborhood of an evolutionarily stable state. In order to clarify this point, a concept of evolutionary stability of genetic populations will be developed. Thus, in a second part of the paper it becomes possible to explain both, the reasons for the specific dynamics of standard GA learning models and the different kind of dynamics of GA learning models, which use extensions to the standard GA.

1 Introduction

Genetic Algorithms (GAs) have been frequently used in economics to characterize a well defined form of social learning.¹ They have been applied to mainstream economics problems and mathematically analyzed as to their specific dynamic and stochastic properties.² But, although widely seen as conducting a rather evolutionary economic line of thought, up to now there is no piece of work explicitly focusing on what it is that makes GA learning an evolutionary kind of behavior.

This paper aims to clarify the scientific advantage of viewing genetic algorithms as evolutionary processes. Evolutionary game theory delivers some very useful tools, which can help explaining why economic GAs behave the way they do. In more detail, it can be found out what are the reasons for the special, non-converging dynamics of the standard GA, and — more than this — it can be explained why certain changes to the GA, like the introduction of the election operator³, can change the GA-dynamics to an as dramatic extent as they do.

As genetic algorithms have been well introduced into economic research, this paper will not explicitly review the specific structure and working principles of GAs. The reader will be provided to have a basic notion of genetic algorithms, which can be gained from e.g. Goldberg (1989) or Mitchell (1996). The kind of genetic algorithms found in these books will be called 'basic' or 'standard' genetic algorithm. In this paper, the standard GA will mainly be dealt with. Nevertheless, the analysis that will be carried out will give rise to the opportunity of analysing more difficult variants of genetic algorithms. Thus, in a further part of this paper, even enhanced or augmented genetic algorithms will be briefly dealt with, clarifying the reasons why some of these variants behave differently from the standard GA.

This paper will face the question what is evolutionary in GA learning. First, it will be shown that there is a close connection between evolutionary game theory and genetic algorithm learning. Using this notion it will be pointed out that genetic evolving populations can be interpreted as a series of near Nash equilibria. Then, in a second step, the well-known concept of evolutionary stability will be transferred to the field of GA research. In a third step, it will be shown that under the regime of the market genetic algorithm learning leads to a kind of evolutionary dynamics, which can be described as economic progress rather than just economic change. A further part of the paper makes clear why some frequently used modifications to the basic GA result in kinds of dynamics which are very different from the basic dynamics. At last, the paper ends with a summary.

2 The Standard Genetic Algorithm

As there is a growingly large number of different variants of Genetic Algorithms in economic research, this paper will mainly deal with the most basic GA, the standard GA, which is

¹Some frequently cited papers are Andreoni and Miller (1995), Arifovic (1994, 95, 96), Axelrod (1987), Birchenhall (1995), and Bullard and Duffy (1998).

²Dawid (1994), Riechmann (1998a).

³See Arifovic (1994).

described in detail by Goldberg (1989). The standard GA is the GA all other variant GAs are derived from.

The standard GA is a stochastic process which repeatedly turns one population of bit strings into another. These bit strings are called genetic individuals. In economic research they normally stand for some economic strategy or routine in the sense of Nelson and Winter (1982). These routines are said to be used by economic agents.⁴

Each repeated ‘turn’ of the genetic algorithm essentially consists of two kinds of stochastic processes, which are variety generating and variety restricting processes. Variety generating processes are processes in that new economic strategies are developed by the economic agents. In the standard GA these processes are reproduction, which is interpreted as learning by imitation, crossover, which is interpreted as learning by communication, and mutation, which is interpreted as learning by experiment. All these processes, or genetic operators, take some old economic strategies and use them to find new ones, thus enhancing the variety of strategies within the current population. In the standard GA, there is one variety restricting process, which is the genetic operator of selection. Selection decreases the number of different economic strategies. It first evaluates the economic success of each strategy, thus often being interpreted as playing the role of the market as an information revealing device.⁵ Then it selects strategies to be part of the next population. The selection operator of the standard GA does so by applying a kind of biased stochastic process: Selection for the next generation is done by repeatedly drawing with replacement strategies from the pool of the old population to be taken over into the next one. The chance of each strategy to be drawn is equal to its relative fitness, which is the ratio of its market success to the sum of the market success of all strategies in the population. Thus, the number of different strategies within a population is reduced again.

3 Genetic Algorithms as Evolutionary Processes

Close relationships between economic learning models and models of evolutionary theory have been recognized before. Marimon (1993) gives a clear notion of the similarities of learning models on the one hand and evolutionary processes on the other. As genetic algorithms, too, have been broadly interpreted as models of economic learning⁶, this section will show that they can also be regarded as evolutionary processes.

At a first glance, it is the structure of genetic algorithms and evolutionary models that seems to suggest a close relationship between GAs and evolutionary economic theory: Both face the central structure of a population of economic agents interacting within some well defined economic environment and aiming to optimize individual behavior.

As the aim of this paper is to describe genetic algorithms as evolutionary processes, the first question to be answered is the question, if GAs are evolutionary processes at all. In the follow-

⁴It is important to clarify the following point: A genetic individual is not interpreted as an economic agent, but as an economic strategy *used by* an economic agent. This interpretation allows for several agents employing the same strategy.

⁵Note, that this has already been described by Hayek (1978).

⁶For such an interpretation see e.g. Dawid (1996).

ing, it will be argued that GAs are a specific form of evolutionary processes, i.e. evolutionary games.

Friedman (1998, p. 16) gives three characteristics for an evolutionary game:

‘By an evolutionary game, I mean any formal model of strategic interaction over time in which (a) higher payoff strategies tend over time to displace lower payoff strategies, but (b) there is some inertia, and (c) players do not systematically attempt to influence other players’ future actions.’

Prior to checking these three points, it is important to notice that economic GAs are in fact models of strategic interactions. In the interpretation as models of social learning, GAs deal with a number of economic agents, each trying to find a behavioral strategy which, due to her surrounding, gives her the best payoff possible. GAs are models of *social* learning, which in fact is a way of learning by interaction.⁷ Thus, it can be recognized, that GAs are in fact models of ‘strategic interaction’.

Moreover, GAs are dynamic processes which reproductively prefer higher payoff strategies over lower payoff ones. It has been shown that in the standard GA the probability of a strategy i to be reproduced from its current population \bar{n} into next period’s population, $P(i|\bar{n})$, only depends on its relative fitness $R(i|\bar{n})$, which is the strategy’s payoff or market success relative to the aggregate payoff of the population \bar{n} .⁸ Higher relative fitness leads to a higher reproduction probability:

$$\frac{dP(i|\bar{n})}{dR(i|\bar{n})} > 0 \quad (1)$$

Thus, Friedman’s condition (a) is fulfilled.

Secondly, according to Friedman, inertia means that changes in behavior do not take place too abruptly. As mutation, or learning by experiment, is the source of the most abrupt changes in GA learning, it should be proved that small changes by mutation are more likely than big ones. For the standard GA, using binary representation of genetic individuals, the mutation operator is quite simple. Mutation randomly alters (‘flips’) single bits of the bit string by which an economic strategy (i.e. a genetic individual) is coded. Each bit of the string has a small probability μ to be changed. μ , which is called the mutation probability, is the same for every bit within every genetic individual of every population. Figure 1 shows an example of the mutation operator. The fact that small changes by mutation are more likely than big ones can be shown as follows: The probability of an economic strategy i to be turned into strategy j by mutation, $P_m(i, j)$, depends on the length of the genetic individuals’ bit strings, L , the mutation probability μ , and the number of bits that have to be flipped in order to turn i into j , which is called the Hamming distance between i and j , $H(i, j)$:

$$P_m(i, j) = \mu^{H(i, j)} (1 - \mu)^{L - H(i, j)} \quad (2)$$

⁷There are, on the contrary, ways of individual, i.e. non-social learning as e.g. statistical forms of learning or neural network learning.

⁸For a more precise description of this, refer to Riechmann (1998a).

```

0 0 1 0 0 0 0 0 0 0 1 0 0 0 1
      |
0 0 1 0 0 0 0 0 1 0 0 1 0 0 0 1

```

Figure 1: Mutation (example)

Simple differentiation gives

$$\frac{\partial P_m(i|j)}{\partial H(i,j)} = \mu^{H(i,j)} (1-\mu)^{L-H(i,j)} [\ln \mu - \ln(1-\mu)] \begin{cases} < 0 & \text{for } \mu < \frac{1}{2} \\ = 0 & \text{for } \mu = \frac{1}{2} \\ > 0 & \text{for } \mu > \frac{1}{2} \end{cases} \quad (3)$$

For the normal parameter value of μ ,⁹ this means the obvious: Small changes in strategy are more likely than big changes.¹⁰ Thus, it becomes evident that GA learning processes are processes which contain some inertia.

Friedman's third point ('players do not systematically attempt to influence other players' future actions') can be proved more verbally. The agents that are modelled by an economic GA have very restricted knowledge. By the time an economic agent forms her latest economic strategy she does not know anything about the plans of the other agents in her population. All an economic agent in a GA model can do is to try her best to adopt to her neighbors' *past* actions, for the near past is all an economic agent can remember. Taking into account these very limited individual abilities, it is easy to conclude, that there is no room for *systematic* influences on other agents' actions.

From the above it can be concluded, that models of economic GA learning are in fact models which can be interpreted as evolutionary games as well.

4 Populations as Near Nash Equilibria

The main structure in genetic algorithm learning models is the genetic population. It can be noticed that a population is nothing more than a distribution of different economic or behavioral strategies.¹¹ This is true for genetic populations as well as for populations in the game theoretic meaning of the word. Thus, it can be said that a genetic population *is* a game theoretic

⁹Normally, the mutation probability μ ranges somewhere between 1/100 and 1/1000.

¹⁰The result given in (2) deserves two further remarks. First, the fact that for $\mu > 1/2$ big changes are more likely than small ones explains that fact that for relatively large values of μ , GA results seem to become very similar to random walks. Secondly, the result yields an interesting interpretation for the field of economic learning. If mutation is interpreted as learning by experiment, (2) shows that a little experimenting is a good thing to do, while too many experiments will disturb the generation of valuable new economic strategies. If mutation is interpreted as making mistakes in imitation or communication (see e.g. Alchian (1950)), (2) simply means that you should not make too many of those mistakes.

¹¹For an in-depth discussion of this, refer to Davis and Principe (1993) or Riechmann (1998a).

population.¹²

Genetic algorithms are in fact describing a repeated economic game. Imagine a genetic algorithm using a population of M genetic individuals with the length of each individual's bit string of L . This GA is able to deal with every economic strategy in S , the set of all available strategies. S has the size $N = |S| = 2^L$. This means that this GA can be interpreted as a repeated symmetric one population M person game with up to N possible pure strategies. But, compared to 'normal' evolutionary games, within most economic GA learning models, the rules are different. Whereas in evolutionary games most of the time a strategy is repeatedly paired with single competing strategies, in genetic algorithm learning, each strategy plays against the whole aggregate rest of the population.¹³ There is no direct opponent to a single strategy. Instead, every economic agent aims to find a strategy $i \in S$ that performs as good as possible relative to its environment, which is completely determined by the current population $\bar{\pi}$ and the objective function $R(\cdot)$.¹⁴

This means that every economic agent i faces problem (4).¹⁵

$$\max_{i \in S} R(i | \bar{\pi}) \quad (4)$$

This directly leads to the concept of Nash equilibria. While a Nash strategy is defined as the best strategy *given the strategies of the competitors*, a Nash strategy is exactly what every economic agent, alias genetic individual, is trying to reach. Thus, as a first step of analysis, a genetic algorithm can be seen as modelling a system of economic agents, each trying to play a Nash strategy. In economic terms this means that every agent tries to coordinate her strategy with the other agents' ones, for this is the best way of maximizing her profit (or utility or payoff or whatever the model wants the agent to maximize). A genetic population can be interpreted as a population of agents, each trying to play Nash.

5 Evolutionary Stability of Genetic Populations

While a genetic population represents a primarily static concept, learning processes are of course genuinely dynamic processes. Thus, in order to analyze GA learning as an evolutionary learning process, the dynamics and concepts of stability have to be analyzed.¹⁶

¹²See Witt (1993a, pp. 4–6) and his description of a 'shift to population thinking' taking place in evolutionary economics.

¹³While this notion is true for most of the economic GA models, for some it is not, including Axelrod (1987) and Andreoni and Miller (1995).

¹⁴The exact mathematical formulation can be found in equation Riechmann (1998a, equation (7)), which, in game theoretic terms, gives the payoff to agent i playing against the rest of population $\bar{\pi}$. It should be noted, that in games a player's payoff depends on his action and the action of every opponent, so that the best formulation of fitness or payoff is $R(i|\bar{\pi})$.

¹⁵Although (4) looks a bit complicated, even compared to most of the mainstream economic models, it is in fact remarkably simple. All it says is 'Do the best you can with respect to your neighborhood!'

¹⁶Replicator dynamics (see e.g. Weibull (1995) or Hofbauer and Sigmund (1988, 98)), which have often been used to characterize evolutionary dynamics, seem to be unsuited for some economic problems. (Mailath (1992, p. 286) even suggests that 'There is nothing in economics to justify replicator dynamics'.) Applied to the analysis of GA learning, replicator dynamics, not directly accounting for stochasticity, are simply not precise enough to cover the whole GA learning process.

This paper will make use of the concept of evolutionary stability, especially the notion of evolutionarily stable strategies or evolutionarily stable states (ESS).¹⁷ In short, a strategy is evolutionarily stable if, relative to its population, it performs better than any new, ‘invading’ strategy. Though widely used in economic dynamics, the concept of ESS has two weaknesses which make this concept seem to be of only limited suitability for the analysis of genetic algorithms. The first weakness lies in the fact that the concept of ESS is based on symmetric two person games only. As mentioned above, this is not the form of games usually played in GA learning. Most GAs have each genetic individual playing against the aggregate rest of the population. Secondly, there is no explicit formulation of the selection process underlying the concept of evolutionary stability. ESS are based on the notion that invading ‘mutant’ strategies are somehow rejected or eliminated from the population. It is not clear how this rejection will be carried out. Genetic algorithms, in contrast, present a clear concept of rejection: Every strategy will be exposed to a test, which is best described as a one–against–the–rest game. Then it will be reproduced or rejected with a probability depending on its performance (i.e. market performance) in the game. GA reproduction or rejection has two main features, it selects due to performance and it selects due to probability, which means that a bad strategy will be rejected almost surely but not with probability one.

Thus, a refined concept of evolutionary stability for genetic algorithms is presented. An attempt to set up a concept of evolutionary stability for genetic algorithms which is keeping the spirit of the ESS is the following: A genetic population is evolutionarily stable if the process of the genetic algorithm rejects an invasion by one or more strategies from the genetic population. Invasion itself can either take the form of a totally new strategy entering the population or it can simply mean a change in the frequency of the strategies already contained within the population. Thus, a clearer definition of an evolutionarily stable population might be: A population is evolutionarily stable if it is resistant against changes in its composition.

More formally, a genetic population \bar{n} will be called *evolutionarily superior* to population \bar{m} , (denoted as $\bar{n} \overset{\text{es}}{>} \bar{m}$) if it exhibits two characteristics¹⁸:

- a) Every strategy i contained within population \bar{n} has at least the same fitness in the basic population \bar{n} as it has in the invaded population \bar{m} , while at least one strategy has even more fitness in \bar{n} than in \bar{m} .
- b) The invading strategies $k \in \{\bar{m} \setminus \bar{n}\}$ are the worst performing strategies contained in \bar{m} , so that they will be most surely rejected.

In mathematical terms, a genetic population \bar{n} is evolutionarily superior to \bar{m} , if

$$R(i|\bar{n}) \geq R(i|\bar{m}) \quad \forall i \in \bar{n} \tag{5}$$

$$\wedge \exists j \text{ with } R(j|\bar{n}) > R(j|\bar{m}) \tag{6}$$

$$\wedge R(k|\bar{m}) < R(i|\bar{m}) \quad \forall i \in \bar{n}; \forall k \in \{\bar{m} \setminus \bar{n}\} \tag{7}$$

¹⁷See Maynard Smith (1982), Hofbauer and Sigmund (1988, 98), Samuelson (1997), Weibull (1995), Marks (1992), or Mailath (1992) (to mention just a few of various pieces of work on this topic).

¹⁸Note that the following characterizes a kind of weak dominance concept.

For full validity, a further remark is necessary, even if it is a little beyond the scope of this paper: Within genetic algorithms, invading strategies can only result from reproduction ('imitation'), crossover ('communication') or mutation ('experiment') within the population itself. This means that the final outcome of GAs without mutation (i.e. processes with learning by imitation and communication only), which are uniform populations, may have other populations being superior to them, but — without mutation — better populations simply cannot arise.¹⁹

Note that (5) to (7) induce a partial ordering on the set of genetic populations S' . A stable population in the concept of (5) to (7) is a population that is superior to every other population: n is an evolutionarily stable population, if

$$\bar{n} \stackrel{\text{es}}{>} \bar{m} \quad \forall \bar{m} \in \{S' \setminus \bar{n}\} \quad (8)$$

Condition (8) is in fact a generalization of the concept of evolutionary stability.²⁰

Due to the fact that genetic algorithm selection is a probabilistic rather than a deterministic process, invading strategies, even in a evolutionarily stable population, may not be rejected within a single round of the algorithm. It can only be stated that the invader will be driven out of the population within finite time. That is to say: If a genetic population is evolutionarily stable, it will recover from an invasion within a finite number of steps of the GA, which means that *in the long run* the population will not lastingly be changed. Nevertheless, once an evolutionarily stable population is invaded, there may appear a number of evolutionarily inferior populations within the next few rounds of the GA. These populations represent transitory states of the process of rejecting the invader. Riechmann (1998a) shows that there is in fact more than one population that will occur in the long run. More precisely, Markov chain analysis of genetic algorithms shows that due to the special dynamic properties of genetic algorithms there is a whole distribution of genetic populations that will repeatedly be reached during the GA process.

6 Evolutionary Dynamics

As a consequence of what has been developed in the preceding parts, dynamics of genetic algorithms can be characterized in a more evolutionary manner. First of all it can be noticed that every population describes a game theoretic outset which is a near Nash equilibrium. The genetic algorithm as a process of turning one population into another can be viewed as at least an approximation of the moving Nash equilibria process. More than this, turning to the criterion of evolutionary superiority ((5) to (8)), the GA always selects in favour of the superior population. This notion can be used to characterize genetic learning dynamics: The stochastic process GA continuously discards populations in favour of better ones in the sense of criterion (8). This only describes the direction of the process, not the exact path that is taken in time. In fact, due to the stochastic properties of genetic algorithms, the exact path of the process

¹⁹See Riechmann (1998a) for the restrictions different learning techniques put on the set of available strategies.

²⁰Note e.g. the similarity to the Weibull's (1995), pp. 36 definition.

highly depends on the starting point, i.e. the composition of the very first genetic population. And although the path up to an evolutionarily stable equilibrium may differ, Markov chain theory shows that it will be reached, and specifically that it will be reached irrespectively of the starting conditions. There may be path dependence, lock-ins, or related phenomena, but in the case of genetic algorithm learning these will only be of temporary nature. In the long run, genetic algorithm theory promises, the 'best' state will be reached.²¹

Knowing the special form of the dynamic process of the GA and the direction in which this process will lead, a few more words can be said about the role of heterogeneity for the dynamics.

It seems important to notice the way economic changes take place. Starting with some population, genetic operators (i.e. learning mechanisms) cause changes in the population. New types of behavior are tested. The test is performed by exposing the strategies to the market. The market reveals the quality of each tested strategy relative to all the other strategies within the population. Then selection lets economic agents give up poorly performing strategies and adopt better ones (imitation) or even create new ones by communication (crossover) or experimentation (mutation). After that, again, strategies are tested and evaluated by the market, by that way coordinating the agents' strategies.

There are in fact two crucial points to this repeated process: First, it is the diversity of strategies that drives economic change, i.e. the succession of populations constantly altering their composition. Under the regime of genetic algorithm learning, this change in individual as well as in social behavior heavily (while not entirely) relies on learning by imitation and learning by communication. As was pointed out in greater detail in Riechmann (1998a), these kinds of learning can only take place within heterogeneous populations. Thus, in a way, it can be said that it is heterogeneity that is the main force behind economic change.

The second crucial point to the process of genetic algorithm learning is the role of selection, which can be interpreted as the role of the market. While the act of learning will be enough to achieve economic *change*, economic *development* can only be reached by the cooperation of learning and selection. In order to turn the succession of different populations into the succession of constantly improving populations (in the sense of evolutionary superiority), a device is needed that makes it possible to distinguish successful strategies from less successful ones. Having at hand such a device, it is possible to decide which strategies shall live and grow and which ones shall die. This device is the market in economics as it is the selection operator within genetic algorithms. It is the market and only the market that turns economic change into economic development.²²

²¹This may be regarded as a weakness of the concept of genetic algorithm learning, as it neglects the possibility of modelling path dependence or lock-ins. So it may be worthwhile to mention two further points, which are mainly beyond the scope of this paper. First, depending on the underlying (economic) problem, some GAs spend long times supporting populations which are not evolutionarily stable. Some keywords pointing to this topic are 'deceptiveness' of genetic algorithms and the problem of 'premature convergence'. Secondly, the lack of ability to model lasting lock-ins or path dependence applies to the basic genetic algorithm. There are variations of genetic algorithms which are capable of modelling these phenomena. One keyword pointing into this direction of research may be 'niching mechanisms'. Again, a good starting point for more descriptions of all of the special cases and variants of GAs is Goldberg (1989).

²²This reflects a rather classical economic thought, given, e.g., in Hayek (1969) (usually quoted as Hayek (1978)).

Summarizing, under the regime of the market, evolutionary dynamics of genetic algorithm learning is mainly driven by two forces: Heterogeneity, which constantly induces behavioral (and by that, economic) change, and the market as a coordination device, revealing information about the quality of each type of behavior and ruling out poorly performing strategies, thus turning economic change into economic development.

Finally, looking at genetic algorithm learning from an evolutionary point of view, one more point has to be added. It has been shown that, as long as possible, genetic algorithm learning and market selection improve individual and with that social behavior. Yet, once an evolutionarily stable state of behavior has been reached, there certainly is no room for further improvement. But, due to the special structure of genetic algorithms, this does not mean that in this state economic agents stop changing their behavior. Learning, or what has above been called change, still continues and will not cease to continue. Still, there will appear new ways of individual behavior within a population. Now it is the role of the market (i.e. selection) to drive these strategies out of the population again. Due to the probabilistic nature of the GA, this process may take more than one period, thus producing one or even more transitory populations until the evolutionarily stable population is regained. To put it in different words: Even after an evolutionarily stable state is reached, evolutionary stability is continuously challenged by new strategies. While in the first phase of the GA learning process some of these new strategies are integrated into the population, in the second phase all of the invaders will be replaced again. So there is an ongoing near equilibrium movement resulting from the continuous rejection of invading strategies.

In fact, genetic algorithm learning leads to an

‘interplay of coordinating tendencies arising from competitive adaptations in the markets and de-coordinating tendencies caused by the introduction of novelty’

(Witt (1993b, p. xix)), which has often been regarded as a key feature of evolutionarily economic analysis of the market.²³

Long run dynamics of GA learning processes have in mathematical terms been characterized as a state of Ljapunov stability (Riechmann (1998a)). With the help of evolutionary game theory, a clear economic reason can be found, why this state of Ljapunov stability shows up: It is a process of near equilibrium dynamics, caused by the continuously ongoing challenge of the ESS by new strategies and the rejection of these strategies that prevents social behavior from total convergence but still keeps it near enough to the equilibrium.

7 Modified Genetic Operators and Their Impact on Stability

Within this paper, only the most basic type of genetic algorithms has been looked at. In economic research, various forms of GAs are used which employ modifications of the operators described in this paper. Some modifications of genetic operators have a major impact on the

²³See Witt (1985).

dynamics and accordingly on the stability properties of the genetic algorithm containing them. Two such modifications will briefly be mentioned.

7.1 Selection

Within this paper the standard form of the selection operator has been analyzed. While this ‘biased roulette wheel’ selection used in the standard genetic algorithm leads to the reported results, there are different types of selection which lead to algorithms with different behavior.²⁴

Above all, there is a group of elitist selection operators, including the selection within evolution strategies²⁵ and Arifovic’s (1994) election operator.²⁶ Elitist selection ensures that at least the first best genetic individual of a population will become a member of the next generation’s population. These differing selection operators show up a much stronger tendency of leading to strict asymptotic convergence and uniformity of genetic populations.²⁷ This tendency can easily be explained. In contrast to roulette wheel selection, elitist selection ensures that invading strategies which turn out to be the worst strategies throughout the population will be replaced at once. This means that there will be no room for transitory populations. Bad strategies, i.e. strategies obeying condition (7), will be ruled out before they can even enter a population. This certainly leads to asymptotic behavioral stability.

Which selection operator to choose for a genetic algorithm in an economic model heavily depends on the economic interpretation of the operators and on its relevance to the problem to be modeled. This interpretation, in turn, mainly depends on the role, the author of the model wants to assign to e.g. chance, mass phenomena, network externalities and related topics. The problem concentrates upon the question whether the best behavior in a certain period will inevitably find its way into next period’s pool of behavioral strategies (elitist selection) or if there will be any forces that can prevent this (roulette wheel selection).

7.2 Mutation

The analysis performed above shows that mutation is a strong force behind economic change. Yet, used as a metaphor of economic learning by experiment, mutation in its simplest form may be seen as underestimating economic agents’ rational capacity. Why should agents not notice if their repeated experiments cease to gain improvements? A modified mutation operator could be thought of, endogenizing each agents propensity to experiment. A modification of this type, based on earlier papers by Bäck (1992a, b) and Bäck and Schütz (1996), has been analyzed by Riechmann (1998b), who finds this change in mutation to smooth but not to totally remove the resulting near equilibrium dynamics of the GA. The reasons for this finding can be found in the fact that once a relatively good state is reached, mutation probability is reduced. Learning by experiment decreases if there is not much left to learn. Thus, there are less invading strategies producing less transitory populations which leads to a slow down in economic fluctuations.

²⁴An overview of various selection schemes can be gained from Goldberg and Deb (1991).

²⁵For a survey, refer to Bäck, Hoffmeister, and Schwefel (1991).

²⁶For an interpretation and an extension of the election operator, see Franke (1997).

²⁷A convergence analysis for genetic algorithms with elitist selection has been carried out by Rudolph (1994). In his work, Rudolph proves that genetic algorithms with elitist selection *will* converge to a uniform population.

Again, which kind of mutation operator to choose mainly depends on the economic interpretation which should be applied to it.

8 Summary

Learning by genetic algorithms is a specific form of a repeated evolutionary game. This fact gives you at hand the whole range of analytical tools evolutionary game theory offers for the analysis of dynamic processes.

This paper proves that GA learning in fact is an evolutionary game. It uses the notion of Nash equilibria and a transferred concept of evolutionary stability to describe in detail the dynamics of genetic algorithm learning both in its standard form and in some of its variants. Though this is just the beginning of some more pieces of work still to be done, the results are rather enlightening and help explain why GA learning works the way it apparently does.

References

- Alchian, Armen A. (1950). Uncertainty, Evolution, and Economic Theory. *Journal of Political Economy*, **58**, 211–221.
- Andreoni, James and Miller, John H. (1995). Auctions with Artificial Adaptive Agents. *Games and Economic Behavior*, **10**, 39–64.
- Arifovic, Jasmina (1994). Genetic Algorithm Learning and the Cobweb–Model. *Journal of Economic Dynamics and Control*, **18**, 3–28.
- Arifovic, Jasmina (1995). Genetic Algorithms and Inflationary Economies. *Journal of Monetary Economics*, **36**, 219–243.
- Arifovic, Jasmina (1996). The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies. *Journal of Political Economy*, **104**, 510–541.
- Axelrod, Robert (1987). The Evolution of Strategies in the Iterated Prisoner’s Dilemma. In: *Genetic Algorithms and Simulated Annealing*, edited by Davis, Lawrence, pp. 32–41. Pitman, London.
- Bäck, Thomas (1992a). The Interaction of Mutation Rate, Selection, and Self–Adaption Within a Genetic Algorithm. In: *Parallel Problem Solving from Nature 2*, edited by Männer, Reinhard and Manderick, Bernhard. Elsevier Science, Amsterdam.
- Bäck, Thomas (1992b). Self–Adaption in Genetic Algorithms. In: *Proceedings of the First European Conference on Artificial Life*, Cambridge, Mass., London, England. MIT Press.
- Bäck, Thomas, Hoffmeister, Frank, and Schwefel, Hans-Paul (1991). A Survey of Evolution Strategies. In: *Proceedings of the 4th International Conference on Genetic Algorithms*, edited by Belew, Richard K. and Booker, Lashon B., pp. 2–9, San Mateo, California. Morgan Kaufmann.

- Bäck, Thomas and Schütz, Martin (1996). Intelligent Mutation Rate Control in Canonical Genetic Algorithms. In: *Foundation of Intelligent Systems 9th International Symposium, ISMIS 96*, edited by Ras, W. and Michalewicz, M., pp. 158–167, Berlin, Heidelberg, New York. Springer.
- Birchenhall, Chris (1995). Modular Technical Change and Genetic Algorithms. *Computational Economics*, **8**, 233–253.
- Bullard, James and Duffy, John (1998). A Model of Learning and Emulation with Artificial Adaptive Agents. *Journal of Economic Dynamics and Control*, **22**, 179–207.
- Davis, Thomas E. and Principe, Jose C. (1993). A Markov Chain Framework for the Simple Genetic Algorithm. *Evolutionary Computation*, **1**, 269–288.
- Dawid, Herbert (1994). A Markov Chain Analysis of Genetic Algorithms with a State Dependent Fitness Function. *Complex Systems*, **8**, 407–417.
- Dawid, Herbert (1996). *Adaptive Learning by Genetic Algorithms*. Springer, Berlin, Heidelberg, New York.
- Franke, Reiner (1997). Behavioural Heterogeneity and Genetic Algorithm Learning in the Cobweb Model. Discussion Paper 9, IKSF — Fachbereich 7 — Wirtschaftswissenschaft. Universität Bremen.
- Friedman, Daniel (1998). On Economic Applications of Evolutionary Game Theory. *Journal of Evolutionary Economics*, **8**, 15–43.
- Goldberg, David E. (1989). *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison–Wesley, Reading, Massachusetts.
- Goldberg, David E. and Deb, Kalyanmoy (1991). A Comparative Analysis of Selection Schemes Used in Genetic Algorithms. In: *Foundations of Genetic Algorithms*, pp. 69–93, San Mateo, California. Morgan Kaufmann.
- Hayek, Friedrich A. von (1969). *Freiburger Studien*, chap. 15, Der Wettbewerb als Entdeckungsverfahren, pp. 249–265. J. C. B. Mohr (Paul Siebeck), Tübingen.
- Hayek, Friedrich A. von (1978). *New Studies in Philosophy, Politics, Economics and the History of Ideas*, chap. 12, Competition as a Discovery Process, pp. 179–190. Routledge & Kegan Paul, London.
- Hofbauer, Josef and Sigmund, Karl (1988). *The Theory of Evolution and Dynamical Systems*. Cambridge University Press, Cambridge, UK.
- Hofbauer, Josef and Sigmund, Karl (1998). *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge, U.K.
- Mailath, George J. (1992). Introduction: Symposium on Evolutionary Game Theory. *Journal of Economic Theory*, **57**, 259–277.

- Marimon, Ramon (1993). Adaptive Learning, Evolutionary Dynamics and Equilibrium Selection in Games. *European Economic Review*, **37**, 603–611.
- Marks, R. E. (1992). Breeding Hybrid Strategies: Optimal Behaviour for Oligopolists. *Journal of Evolutionary Economics*, **2**, 17–38.
- Maynard Smith, John (1982). *Evolution and the Theory of Games*. Cambridge University Press, Cambridge, UK.
- Mitchell, Melanie (1996). *An Introduction to Genetic Algorithms*. MIT Press, Cambridge, Mass., London, England.
- Nelson, Richard R. and Winter, Sidney G. (1982). *An Evolutionary Theory of Economic Change*. MIT Press, Cambridge, Mass.
- Riechmann, Thomas (1998a). Learning and Behavioral Stability — An Economic Interpretation of Genetic Algorithms. *Journal of Evolutionary Economics*, forthcoming.
- Riechmann, Thomas (1998b). Learning How To Learn. Towards an Improved Mutation Operator within GA Learning Models. In: *Computation in Economics, Finance and Engineering: Economic Systems*, Cambridge, England. Society for Computational Economics.
- Rudolph, Günter (1994). Convergence Analysis of Canonical Genetic Algorithms. *IEEE Transactions on Neural Networks*, **5**, 96–101.
- Samuelson, Larry (1997). *Evolutionary Games and Equilibrium Selection*. MIT Press Series on Economic Learning and Social Evolution. MIT Press, Cambridge, Mass.
- Weibull, Jörgen W. (1995). *Evolutionary Game Theory*. MIT Press, Cambridge, Mass., London, England.
- Witt, Ulrich (1985). Coordination of Individual Economic Activities as an Evolving Process of Self-Organization. *Economic Appliquée*, **XXXVII**, 569–595.
- Witt, Ulrich (1993a). Evolutionary Economics: Some Principles. In: *Evolution in Markets and Institutions*, edited by Witt, Ulrich, pp. 1–16. Physica-Verlag, Heidelberg.
- Witt, Ulrich (1993b). Introduction. In: *Evolutionary Economics*, edited by Witt, Ulrich, pp. xiii–xxvii. Edward Elgar, Aldershot, England.