J Evol Econ (1998) 8: 297–315

Journal of
Evolutionary
Economics

On economic applications of the genetic algorithm: a model of the cobweb type*

Herbert Dawid[1], Michael Kopel[2]

[1]Department of Management Science, University of Vienna,
Brünnerstrasse 72, A-1210 Wien, Austria (e-mail: herbert.dawid@univie.ac.at)
[2]Department of Managerial Economics and Industrial Organization,
Vienna University of Technology, Wien, Austria

Abstract. This paper explores the idea of using artificial adaptive agents in economic theory. In particular, we use Genetic Algorithms (GAs) to model the learning behavior of a population of adaptive and boundedly rational agents interacting in an economic system. We analyze the behavior of a GA in two versions of a model of the cobweb-type, one in which firms make only quantity choices, and the other one in which firms first decide to exit or to stay in the market, and subsequently decide how much to produce. We present simulations with different coding schemes and interpret the rather surprising differences between the results for different setups by employing the mathematical theory for GAs with state-dependent fitness functions. In particular, we explain the relationship between coding and convergence properties of GAs.

Key words: Cobweb model – Genetic algorithms – Learning – Artificial economic agents

JEL-classification: D83

1 Introduction

The question of learning by adaptive agents has only recently received wide attention in the economics and game-theoretic literature (see Nyarko et al., 1994; Kirman and Salmon, 1995). To introduce learning into the analysis of

---

* This research was supported by the Austrian Science Foundation under contract Nos. P9112-SOZ, J1281-SOZ and J01003-SOZ. Helpful comments of an anonymous referee are gratefully acknowledged.
Correspondence to: H. Dawid

economic models raises a host of fundamental questions concerning the role of rationality in economic behavior. What sort of learning mechanism – Bayesian, adaptive, or some boundedly rational rule of thumb – is appropriate? How does the market interaction and the interaction between individuals modify agents' beliefs and actions? How does the population evolve by diminishing the relative number of agents whose beliefs or actions turned out not to be profitable in the past? Does the learning mechanism converge, and if so, to what? Is the limit of the learning process an equilibrium situation in the economic sense?

The plethora of questions should show "the wilderness of possibilities" (Sargent, 1993, p. 23) in which a relaxation of the assumption of rational expectations leads us. Precisely how are we to go about building models populated by agents who in some sense are behaving as we scientists do, and how do we model adaptation and learning processes of these boundedly rational agents? An increasing number of economists seek to answer (all or some of) the questions above by adapting certain techniques from the field of "computational intelligence", where the most widely used techniques are neural networks, genetic algorithms, classifier systems based on genetic algorithms, and cellular automata. Although the field of artificial adaptive agents has by no means reached the goal of building entire artificial beings comparable to humans (see Holland and Miller, 1991), and the underlying behavioristic routine described by these algorithms is still not clear, it seems promising that the techniques developed are well suited to the imitation of human learning in simple models (see Andreoni and Miller, 1995; Arifovic, 1994, 1995, 1996; Arthur, 1991, 1993). Although the applications of AI methods in economics show their potential in principle, there have been very few successful attempts to apply them to economic problems.[1]

In this paper we restrict our attention to genetic algorithms (GAs) developed by Holland (1975) (for an introduction into GAs, see Goldberg, 1989 or Mitchell, 1995). The most valuable features of GAs from the viewpoint of economics are the explicit representation of every individual in a population of heterogenous agents who (might) differ in strategies, the parallel processing of information, competition among alternative rules, selection of those that perform better and the possibility of creating new rules. Recently, GAs have been used to model economic agents in auctions (Andreoni and Miller, 1995), to evolve strategies for the Classical Prisoner's Dilemma (Axelrod, 1987; Mühlenbein, 1991; Ho, 1996; Miller, 1996) and for a Generalized Prisoner's Dilemma (Marks, 1992) and also to analyze the learning behavior in 3×3 Strategic-Form-Games (Dawid and Mehlmann, 1996) or signaling games (Arifovic and Eaton, 1995). Furthermore, GAs have also been used to analyze the behavior of adaptive agents in standard economic models like the cobweb model (Arifovic, 1994), overlapping generations models (Arifovic, 1995, 1996; Dawid, 1996a; see also the experiments of Marimon and Sunder, 1994), or in a simple Kyotaki–Wright

[1] Benaroch (1995) notes that, "[A] major reason for this is a failure to realize that most AI methods can be useful only once they have been adapted or even tailored to economic problems."(p. 601).

model of monies (Marimon et al., 1990). Using data of a regional U.S. coffee market, Midgley et al. (1997) tried to breed robust strategies, and demonstrated that they outperform the historical actions of brand managers in this regional market (for the limitations of their results and the problems of implementation the authors faced, we refer the reader to the original manuscript).

Here we use GAs to model the learning behavior of a population of adaptive economic agents in a cobweb-type model. We present simulations with different coding schemes[2] and interpret the rather surprising differences between the results for different setups by employing the mathematical theory for GAs with state-dependent fitness function. Abstracting from this example we argue further that theoretical considerations may be very important in predicting the outcome of simulations and designing a GA (see also Benaroch, 1996).

In what follows we introduce the Simple Genetic Algorithm (SGA) and briefly discuss the economic interpretation of the model used (Section 2). We analyze the behavior of a GA in two versions of a model of the cobweb-type, one in which firms make only quantity choices, and the other in which firms first decide to exit or to stay in the market, and subsequently decide how much to produce (Section 3). A discussion and some concluding remarks are given in Section 4.

2 GAs in economic modeling

Typically a genetic algorithm works on a population of binary strings, which correspond to the chromosomes in natural systems. In the standard setup the strings all have the same length, and this length is one of the parameters that have to be set by the researcher. We will denote the length of a string by l. The set of all binary strings with length l is denoted by Ω. The cardinality of this set is given by $|\Omega| = 2^l = r$. Another parameter to be determined is the size of the population. We call this parameter n and always assume that n is even. A state of the population is given by a frequency distribution of these n strings over Ω. Thus, the set of all possible population states is given by $S = \{\phi \in \Delta^r \mid n\phi_k \in \mathbb{N}\, \forall k \in \Omega\}$, where Δ^r is the $r-1$ dimensional simplex and ϕ_k is the frequency of string k in the population state ϕ. Of particular interest are the so-called *uniform states* of the population where all strings are identical. The state where all strings equal k is denoted by the unit vector e_k.There has to be an externally given fitness function, which assigns at every time t some positive fitness value to any string in the population. A key difference between applications of GAs in optimization problems and in economic models is that in the former the fitness of a string typically depends only on the value of the string, whereas in the latter it depends on the entire state of the population. Thus, in economic systems the fitness function is formally given by $f : S \to \mathbb{R}^r_+$, where $f_k(\phi)$ is the fitness of string k in a population in state ϕ. In general,

---

[2] By a coding scheme we denote the mechanism which translates a strategy in the economic model to a binary string in the GA.

the initial population P_0 is generated randomly and the transition from P_t to P_{t+1} is executed by applying several genetic operators. In the simple Genetic Algorithm analyzed in this paper only three standard operators are used:

(i) *Proportional selection*: The selection operator is intended to implement the idea of the "*survival of the fittest*". Basically, the selection operator determines which of the strings in the current population will be allowed to pass their genetic material to the next generation: we say that the selection operator builds up the *mating pool* by selecting n strings from the current population. The standard selection operator (called *proportional selection*) does this by carrying out n random draws with replacement out of P_t. The probability that a given member of the current population is chosen at one certain draw is proportional to the fitness of this string. Accordingly, we expect to have $\frac{n\phi_k f_k(\phi)}{\sum_{j\in\Omega}\phi_j f_j(\phi)}$ strings $k \in \Omega$ in the mating pool, where ϕ is the state of P_t.

(ii) *One-point Crossover*: Inspired by the example of nature, crossover is intended to join the genetic material of strings with a high fitness in order to produce better individuals. The mating pool is split into $n/2$ pairs of strings and the following operator is applied to each pair with the *crossover probability* χ. The value of χ will in general be larger than 0.6, and often $\chi = 1$ is used. With probability $1 - \chi$ no changes are made to both strings, but with probability χ genetic material is exchanged between the two parents. In the simplest case of *one-point crossover*, one crossover point is drawn randomly between 1 and $l - 1$. Afterwards the values of the bits to the right of the crossover point are swapped between the two parents.

(iii) *Mutation*: The mutation operator should allow the GA to find solutions, which contain bit values that are not existent in the current population. The parameter governing this operator is called *mutation probability* and will be denoted by μ. After the application of crossover each bit in any string is inverted (set to 1 if it was 0, and vice versa) with probability μ.

After the application of these operators the new population P_{t+1} is complete. This procedure is repeated until all strings in the population are equal or a prescribed number of iterations has been performed.

We now turn to our interpretation of the procedure described above as an adaptation process of a population of interacting economic individuals. Beforehand we would like to point out that we regard one string as the representation of one individual in the population. Thus each individual is completely characterized by the strategy encoded by this binary string.

For an economic interpretation of (i) consider a population where the individuals are not able to gather enough information to build a reliable estimate about the future development of the economic system, and to determine their optimal reaction to such a development. In such a situation imitation of previously successful strategies is a plausible behavioral assumption (see e.g. Friedman, 1991; Vega-Redondo, 1995, Schlag, 1998). Such an effect is modeled by proportional selection, where strategies which have yielded an above average payoff will, on average, be used by more agents in the next period. The crossover operator described in (ii) introduces an exchange of information, which is rarely incorporated in economic learning models. The exchange of parts of the string may be interpreted as the adoption of certain special details of a competitor's behavior or strategy, which may be due to overt communication or industrial espionage. Thus individuals may influence each others' behavior even if they do not imitate.[3] Finally, the mutation operator introduced in (*iii*) incorporates innovations made by the individuals either by purpose or by chance. Such effects – think for example of the trembling hand (Selten, 1975) – are quite common in economic modelling.[4]

We are aware of the fact that a GA is not a "canonical" representation of the effects described above, but we believe that a sound interpretation of the different effects incorporated in a learning algorithm is necessary to understand simulation as well as analytical results. The fact that GAs may not only be used as an optimization tool but also as a model of an interacting population (an economy) has, for example, been recently pointed out by Goldberg (1995), who claims that "… *much of the mystery of such systems emanates from their innovative nature, and GAs replace the mystery shrouding innovation with a healthy dosage of mechanism.*"

3 The cobweb model

The cobweb model describes temporary equilibrium market prices in a single market with one lag in supply; since production takes time, quantities produced must be decided before a market price is observed. The model was introduced by Leontief (1934), who postulated a linear model of demand and supply in which agents forget all except their most recent experiences (see also Ezekiel, 1938). As is well-known, convergence to an equilibrium price occurs in this model if supply is less elastic than demand. The first dynamic analysis of the classic linear cobweb model when agents have memory was by Nerlove (1958), and was extended by Muth (1961) to an analysis of rational learning with memory. Carlson (1968) concluded that the linear cobweb model was stable when agents use the mean of past prices

[3] It may well be that the substrategy adopted from some other firm's strategy has some undesirable effects in the new firm's strategy. But this seems realistic, since if a firm implements only one part of some other firm's plan of action, this particular part might not work well in the new environment.

[4] Some of the ideas presented above are, of course, not new in economics. As one of the first, Alchian (1950) in his seminal paper explicitly recognized the use of imitation in contrast to optimizing behavior in guiding much economic behavior. He states that "uncertainty provides an excellent reason for imitation of observed success." He further concludes that "Imperfect imitators provide opportunity for innovation, and the survival criterion of the economy determines the successful, possibly because imperfect, imitators." (p. 219). As the economic counterparts of genetic heredity, mutations, and natural selection he relates imitation, innovation and positive profits. These ideas are also incorporated in evolutionary theories in economics, where "Evolutionary models in the social domain involve some processes of imperfect (mistake-ridden) *learning and discovery*, on the one hand, and some *selection mechanism*, on the other." (Dosi and Nelson, 1994, pp. 154; emphasis in original).

as their forecast price. The issue of the convergence of agent's learning about a rational expectations equilibrium has also been addressed in a cobweb model. Bray (1982) and Bray and Savin (1986) demonstrated that, if agents use a least-squares learning procedure, prices in a cobweb model almost surely converge to a rational expectations equilibrium. In contrast to their "ad hoc" learning scheme, Bray and Kreps (1987) demonstrated that, in the context of this model, there is a unique equilibrium with rational learning.

In what follows we will analyze a model of the cobweb type where we include a term for fixed costs or overhead. Such a term may include deduction for depreciation, interest on long term debt, property taxes and executive compensations.[5] We will present simulations of the learning behavior of Genetic Algorithms in such a model and explore the relationship between coding of strategies and convergence properties.

3.1 A model of the cobweb type with fixed costs

There are n firms in a competitive market that are price takers and that produce the same good. Denote the quantity produced by firm i in period t with $y_{i,t}$. Each firm has the same cost of production

$$c(y_{i,t}) = \begin{cases} \alpha + \beta y_{i,t}^2, \quad \alpha, \beta > 0 & y_{i,t} > 0 \\ 0 & y_{i,t} = 0 \end{cases}$$

where α denotes the short term fixed costs of the firm. The decision makers in the firm do not know the price obtaining in the next period when they have to decide how much to produce, but they do have an expected price, $p_{i,t}^e$. Based on this expectation, firm i chooses an output level that makes its expected profit

$$\Pi(p_{i,t}^e, y_{i,t}) = \begin{cases} p_{i,t}^e y_{i,t} - \alpha - \beta y_{i,t}^2 & y_{i,t} > 0 \\ 0 & y_{i,t} = 0 \end{cases}$$

as large as possible. The optimal quantity for firm i is given by

$$y^*(p_{i,t}^e) = \begin{cases} \frac{p_{i,t}^e}{2\beta} & p_{i,t}^e > 2\sqrt{\alpha\beta} \\ \{0, \sqrt{\frac{\alpha}{\beta}}\} & p_{i,t}^e = 2\sqrt{\alpha\beta} \\ 0 & p_{i,t}^e < 2\sqrt{\alpha\beta} \end{cases} \tag{1}$$

The price p_t that clears the market in period t is then determined by the inverse demand function

$$p_t = a - b\sum_{i=1}^{n} y_{i,t} \tag{2}$$

with $a, b > 0$. Note that in order to guarantee that the expected profit is positive at least for some combinations of the expected price and the

[5] Cobweb models with fixed costs were analyzed by Day (1994) and Kopel (1997).

quantity produced, we have to assume that

$$\alpha < \frac{a^2}{4\beta} \quad . \tag{3}$$

Otherwise, there would be no incentive for the firms to enter the market. In a (homogenous) rational expectations equilibrium, firms expectations about the price are equal to the (afterwards) observed equilibrium price, i.e. $p_t = p^e_{i,t}$ for all i, and all firms take the same individual optimal action, $y^*(p^e_{i,t}) = y^*$ for all i for some $y^* \in y^*(p_t)$. According to (1) we have to distinguish two cases. Suppose, that the expected price is greater than $2\sqrt{\alpha\beta}$. Then with (1) and (2) we have

$$a - bny^*(p_t) = 2\beta y^*(p_t)$$

which yields the equilibrium quantity

$$y^* = y^*(p_t) = \frac{a}{2\beta + \gamma} \tag{4}$$

and the equilibrium price

$$p^* = \frac{2a\beta}{2\beta + \gamma} \tag{5}$$

where $\gamma := bn$. The last expression is greater than $2\sqrt{\alpha\beta}$ iff

$$\alpha < \frac{a^2\beta}{(2\beta + \gamma)^2} \quad . \tag{6}$$

On the other hand, if we assume that the expected price is less than $2\sqrt{\alpha\beta}$, we get

$$p^* = a$$

which is less than $2\sqrt{\alpha\beta}$ iff

$$\alpha > \frac{a^2}{4\beta}$$

This, however, contradicts (3).

Hence, in our model two situations might occur, depending on the level of fixed costs: If (6) holds, there exists a unique rational expectations equilibrium with an equilibrium price given by (5) and an equilibrium quantity given by (4). Otherwise, no homogeneous rational expectations equilibrium exists in the model.

3.2 Pure quantity decisions

We simulate the evolving behavior of the system described above, by representing each firm by a binary string. Every binary string b_i in the population P_t encodes a real number in $[0, \frac{a}{\gamma}]$, namely the quantity the firm i decides to produce at time t. We encoded the quantity decisions of the firms as follows: let $k \in \Omega$ be the binary string representing the quantity decision of some firm in the market. The output quantity of this firm is

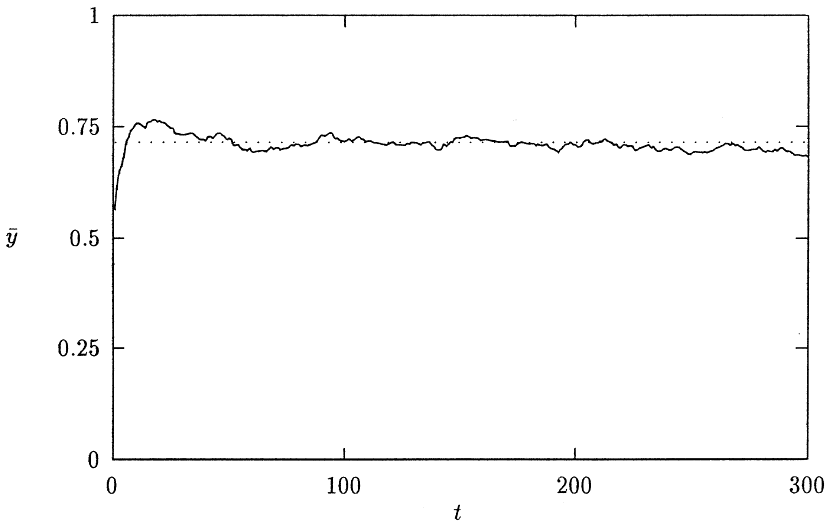

Fig. 1. The average production quantity in P_t for $a = 5,\ \alpha = 0.25$ and $\beta = 1,\ \gamma = 1$. The dotted line describes the rational expectations equilibrium ($n = 1000,\ \chi = 1, \mu = 0.001$)

given by $y(k) = \frac{ax}{\gamma}$, where $x = \sum_{i=1}^{l} 2^{-i}k(i)$ and $k(i) \in \{0, 1\}$ is the value of the i-th bit in string k. The price at time t is calculated using (2). Note that the profit of a firm may become negative for too low prices. The dynamics of the price adaptation of the firms in this market is modeled by a simple genetic algorithm as described in Section 2. To rule out negative fitness values for single strings we use the following scaled profit as the fitness function

$$f_k(\phi) = \Pi(p(\phi), y(k)) + \alpha + \beta\left(\frac{a}{\gamma}\right)^2 , \tag{7}$$

where $p(\phi)$ denotes the price which emerges from (2), if the whole population is in state $\phi \in S$.

In the first simulation we use the following parameter values:

$$a = 5,\ \ \alpha = 0.25,\ \ \beta = 1,\ \ \gamma = 5\ . \tag{8}$$

It is easy to check that conditions (3) and (6) are satisfied for these parameter values, which implies that a unique homogeneous rational expectations equilibrium exists with an equilibrium output of

$$y^* = y^*(p^*) = \frac{a}{2\beta + \gamma} = \frac{5}{7}\ .$$

The length of the string in our first simulation is $l = 10$. We have chosen a rather large population size of $n = 1000$ and crossover and mutation probabilities of $\chi = 1$ and $\mu = 0.001$ respectively.[6] In Fig. 1 the evolution of

[6] We also carried out simulations with different values of the parameters n, l, μ. The behavior of the GA proved to be quite insensitive with respect to these variations.

the average production quantity in P_t is depicted for a simulation with a SGA. We see a very rapid approach towards the equilibrium value. After about 100 generations the population has found the rational expectations equilibrium, and stays in this equilibrium for the rest of the run. The firms are able to adapt their production decisions in such a way that their output quantities are always optimal after a rather short learning period. The convergence of the GA to the rational expectations equilibrium is not surprising as it has been shown by several authors that equilibria may be found by such type of learning algorithms in relatively simple economic models (see e.g. Arifovic, 1994, 1995, 1996).

Let us now consider a slightly different parameter constellation, where the fixed costs are higher than in our first simulation:

$$a = 5, \quad \alpha = 1, \quad \beta = 1, \quad \gamma = 5 \ . \tag{9}$$

Condition (3) is still satisfied, but (6) no longer holds. This implies that in this model no homogeneous rational expectations equilibrium exists. Looking at the simulation in this model (Fig. 2) we observe, however, that the population of firms represented by the GA behaves exactly the same way as in the previous case.

The average produced quantity converges towards the value y^* in (4) and the price, therefore, towards p^* in (5). Thus, the population evolves towards a state where all firms are active and produce a quantity which is optimal for a *producing* firm given the market price. However, the fixed costs of production are so high that they all make negative profits (the profit of the representative firm at $t = 300$ is -0.4763). Obviously, the firms would be better off by exiting the market with a profit of zero. This situation can be seen as some kind of local equilibrium, where no firm can gain by a *small* unilateral deviation from its current strategy. Of course, this is not an economic equilibrium[7] since a firm can gain by unilaterally stopping its production. Note, however, that the state in which all firms abstain from producing is no economic equilibrium either, since (3) implies that in such a situation the market price would be large enough to allow positive profits for firms which unilaterally start producing again. Given this situation the question arises whether it is a pure coincidence that this local equilibrium emerged as the long run state of our simulations or whether we can characterize such states more systematically. In this context it also seems to be crucial to specify what we mean exactly by a small deviation. In order to answer the questions stated above we rely on the mathematical results derived in Dawid (1994) and Dawid and Hornik (1996).[8]

A necessary condition for some state to be the long run outcome of the learning process of the GA is that this state has to be locally asymptotically

---

[7] By an economic equilibrium we refer to a state of the population where every string encodes an optimal strategy given the behavior of the other members of the population.

[8] The basic mathematical model of a GA in economic applications used in the derivation of these results is given in the Appendix.

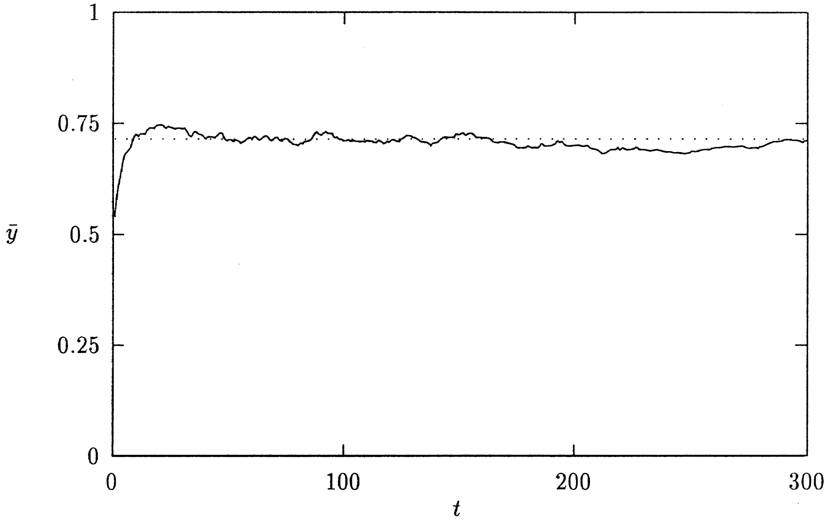

Fig. 2. The average production quantity in P_t for $a = 5$, $\alpha = 1$ and $\beta = 1$, $\gamma = 5$. The dotted line describes the rational expectations equilibrium for the cases where $\alpha < a^2\beta/(2\beta+\gamma)^2$ holds ($n = 1000$, $\chi = 1$, $\mu = 0.001$)

stable with respect to the underlying dynamics[9] of the process. For large population sizes the set of locally asymptotically stable uniform states [10] is characterized by the following proposition.[11]

Proposition 1 *A uniform state e_k is locally asymptotically stable for the dynamics of a GA with $\mu = 0$ and one-point crossover with probability $\chi \in (0, 1]$ if*

$$\frac{d(j \oplus k)}{l-1} > \frac{1}{\chi}\left(1 - \frac{f_k(e_k)}{f_j(e_k)}\right) \tag{10}$$

for all $j \in \Omega$, $j \neq k$, where $d(j \oplus k)$ is the length between the two outer-most bits where j and k differ in value. If there is a string $j \neq k$ where the inequality holds the other way round, e_k is unstable.

Proposition 1 claims that a state representing a population which consists almost only of strings k converges to the uniform state e_k if the strings receiving a higher payoff in the current population differ from k in bits positioned far apart. The intuition behind this result is that strings j, though receiving higher payoffs and thus generating more offsprings, are destroyed

[9] See the Appendix.

[10] We restrict our attention to uniform states since it is shown in Dawid and Hornik (1996) that for small mutation probabilities the process stays "most of the time" near uniform states.

[11] A proof can be found in Dawid (1994).

k	1	0	1	1	0	1	1	0	1	1
0	0	0	0	0	0	0	0	0	0	0

$d(0 \oplus k) = 9$

Fig. 3. The length between the two outmost bits where k and 0 differ in value is $d(0 \oplus k) = 9$

by crossover with a high probability (due to the fact that they differ from k in bits positioned far apart) when paired with a string k.[12]

Let us analyze the simulation results stated above in the light of these theoretical findings. With the parameter values given by (9) the output quantity y^* is given by $y^* \approx 0.7143$. Given our encoding scheme this value is best approximated by the binary string $k = 1011011011$ corresponding to $y \approx 0.71387$. In a state where all strings in the population equal k, the fitness of this string is given by $f_k(e_k) = 1.51$. On the other hand, the production quantity $y = 0$ would yield a profit of zero and due to (7), the fitness of the string encoding this output decision (represented by $j = 0 = 0000000000$) is $f_0(e_k) = \alpha + \beta\left(\frac{a}{\gamma}\right)^2 = 2$. Considering the value of $d(0 \oplus k)$ we realize that these two strings differ as well in their first as in their last bit position. Accordingly, the length between the two outer-most bits where 0 and k differ in value is $d(0 \oplus k) = 9$ (see Fig. 3). Inserting all these values into the stability criterion of proposition 1 yields the result that e_k is locally asymptotically stable since $\frac{9}{9} > \left(1 - \frac{f_k(e_k)}{f_0(e_k)}\right) = 0.245$. This explains the outcome of our simulations also from a theoretical point of view. Furthermore, these arguments show that the expression $d(j \oplus k)$ is in this context a suitable measure for the deviation of the output decision represented by string j from the output decision encoded by k. Although the theoretical results guarantee only local stability of e_k, it turnes out that in simulations where the initial population is initiated with $y = 0$ for all strings, the state trajectory rapidly approaches e_k, hence giving strong evidence for global stability.

The stability of the uniform state e_k may also be derived directly. Assume that in a population of individuals producing $y = 0.71387$ one firm (e.g. by chance) decides to stop producing. As pointed out above this firm would be more successful than the others. We could expect that the strategy to stop producing would be imitated by about $\frac{f_0(e_k)}{f_k(e_k)} = 1.33$ individuals in the population. By contrast, any communication of an inactive firm with some other firm which uses $y = 0.71387$ leads the idle firm to start producing again.[13] Since the communication effect outweighs the imitation effect if the probability of communication (given by χ) is large enough, the number of individuals who use the optimal strategy decreases and the population again converges towards the local equilibrium state e_k.

[12] Note that the left hand side of (10) gives the probability that none of the two offsprings of the parents j and k equals string j if one-point crossover is used.

[13] Formally this is represented by $\frac{d(0 \oplus k)}{l-1} = 1$.

$\hat{k}$	1	1	0	1	1	0	1	1	0	1	1
$\hat{0}$	0	1	0	1	1	0	1	1	0	1	1

$d(\hat{0} \oplus \hat{k}) = 0$

Fig. 4. Change in the coding scheme by enhancing the string by an additional bit. The distance between the resulting strings $\widehat{0}$ and $\widehat{k}$ is $d(\widehat{0} \oplus \widehat{k}) = 0$, and hence the uniform state $e_{\hat{k}}$ is unstable

3.3 Entry and exit decisions

As illustrated above, our simulation result (which shows that the learning algorithm converges to a non-equilibrium state) can be explained from a theoretical standpoint as well as intuitively. We will use these insights to analyze the correspondence between the coding scheme and the learning behavior of the GA.

If we want to avoid a "lock-in" of the GA in a long run state, we need to ensure that this state is unstable with respect to the dynamics of the GA. In order to do this we again invoke the results presented in Proposition 1. Considering the stability criterion in the proposition, it is obvious that the value of the distance $d(0 \oplus k)$ is of crucial importance for the stability properties of the state e_k. If we change the encoding scheme such that the distance between the string representing the production decision $y = y^*$ and the string encoding $y = 0$ is less than $\frac{l-1}{\chi}\left(1 - \frac{f_k(e_k)}{f_0(e_k)}\right)$, the uniform state where all firms produce $y = y^*$ is unstable. A natural way to achieve this is to enhance the strings by one additional bit. This has the following interpretation. Whenever the value of this bit is 0 the firm exits the market, i.e. it stops producing. On the other hand, if the value equals 1, the firm produces the quantity encoded by the remaining l bits. We now have a large number of strings encoding the decision $y = 0$, and changing from $y = y^*$ to $y = 0$ can be accomplished by changing only one single bit. As Fig. 4 shows, with this change in the coding scheme there is a string $\widehat{0}$ encoding $y = 0$; the distance to the string $\widehat{k}$ representing $y = y^*$ is zero. Accordingly, we expect that the uniform state where all firms produce y^* will no longer be the long run outcome. The change in the encoding of strategies also makes economic sense. With the previous coding scheme the basic decision to produce nothing at all and to avoid thereby the fixed costs can only be made by changing all bits in the string to 0. However, this basic decision changes the whole structure of the incurred costs and it seems more plausible to separate the decision of the firm into two parts. With the changed coding scheme the firm decides, first, whether to produce or not and, second, selects the production quantity only in the case the former decision was affirmative. Thus, we have implemented a separation of the production decision into an exit and entry decision and a quantity determination.

As can be seen in Fig. 5, our theoretical considerations are confirmed by the simulations. Figure 5 shows the evolving behavior of the average production quantity of a GA with the extended coding in the same model as used in figure 2. This time, however, the trajectory does not approach the

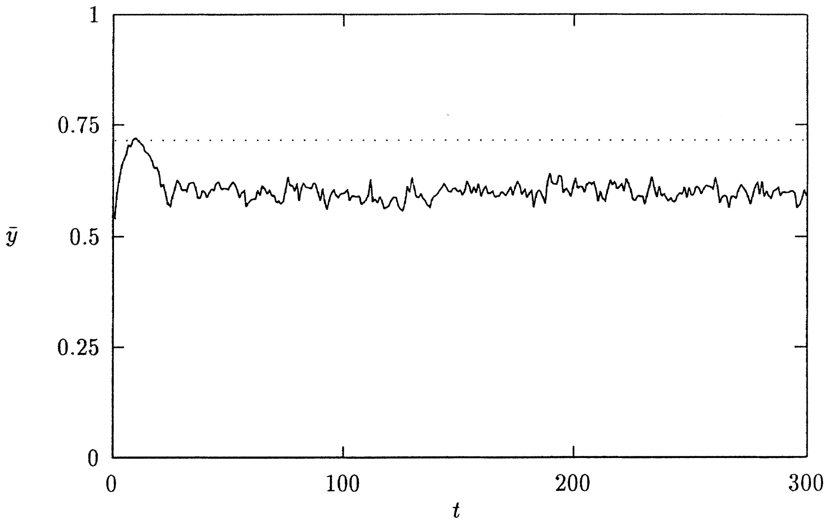

Fig. 5. The average production quantity in P_t for $a = 5,\ \alpha = 1$ and $\beta = 1,\ \gamma = 5$ when the production decision is split into an exit and entry and a quantity decision ($n = 1000,\ \chi = 1, \mu = 0.001$)

value $y = \frac{5}{7}$, but ends up oscillating around $y = 0.6$. The explanation for this behavior is now straightforward. In the beginning of the run $\bar{y}$ is about 0.5, which implies that the price is about $5 - 2.5 = 2.5 > 2\sqrt{\alpha\beta} = 2$. The best decision of the firms is therefore to produce about $y = 1.25$, and due to selection pressure $\bar{y}$ increases above 0.6. Correspondingly the price falls below $p = 2$, which implies that it is now optimal to exit the market. As a switch to this optimal action requires changing one single bit, several firms do this rather quickly. Accordingly, the average output decreases again until the firms producing an amount larger than 0.6 have the highest fitness. Afterwards, these firms are selected and the price falls again. The trajectory then keeps oscillating around $y = 0.6$. Thus, we get a heterogeneous population with a rather high variance, where some of the firms produce nothing at all, whereas others produce a quantity near $y = \sqrt{\alpha/\beta} = 1$.

In fact, the GA has found a heterogeneous rational expectations equilibrium,[14] which describes an industry comprised of active and idle firms. All firms expect the same price $p = 2\sqrt{\alpha\beta}$, but for this expected price the optimal production quantity is not unique. Both $y = 0$ and $y = \sqrt{\alpha/\beta}$ yield the optimal profit of 0. If the fraction of firms producing $y = \sqrt{\alpha/\beta}$ equals exactly $x = \frac{a}{\gamma}\sqrt{\beta/\alpha} - 2\frac{\beta}{\gamma}$ (it is easy to show that $x \in [0, 1]$ if (3) holds and (6) does not hold), the average production quantity is $\bar{y} = x\sqrt{\alpha/\beta}$, and we get from (2)

[14] Hallagan and Joerding (1983) detected a similar situation in an advertising model and called it – motivated by biological models – a polymorphic equilibrium.

$$p = a - \gamma\left(\frac{a}{\gamma}\sqrt{\frac{\beta}{\alpha}} - 2\frac{\beta}{\gamma}\right)\sqrt{\frac{\alpha}{\beta}} = a - a + 2\sqrt{\alpha\beta} = 2\sqrt{\alpha\beta}\ .$$

Thus all firms did expect a price, which subsequently turns out to be the actual market price. We have a heterogeneous rational expectations equilibrium, where all firms hold the same expectations about the price and where the two optimal actions (idle or producing) are chosen in a proportion such that these price expectations are confirmed by the market. This heterogeneous rational expectations equilibrium emerges from the homogeneous equilibrium as the fixed costs pass the value of $\alpha = \frac{a^2\beta}{(2\beta+\gamma)^2}$. Beyond this value, a rational expectations equilibrium can only be adopted if some firms retire from the market and no longer produce. The number of firms producing in the equilibrium decreases for increasing fixed costs until for $\alpha = \frac{a^2}{4\beta}$ no firm will remain in the market, which would mean the demise of the industry.

For our parameter values we get $x = 0.6$, i.e. 400 out of 1000 firms should decide not to produce. In our simulation with $n = 1000$, we observe after 300 generations that there are 378 firms which do not produce. The average profit in the population should be 0. However, due to sampling errors the actual average profit in the simulation shown in Fig. 5 is 0.0262. The oscillations around and the deviations from the theoretical values are higher than in the case of a convergence against a uniform equilibrium[15] since selection may change the ratio between producers and non-producers in addition to the disruptions caused by mutations. Taking this into account, we believe that the approximation of the heterogeneous equilibrium state is quite satisfactory.

We have also carried out further simulation with more than one additional bit. It turnes out that it is not the number of bits which distinguishes the strings in a uniform population from the best reply to it which is crucial for the stability properties. Again – as implied by Proposition 1 – it is the distance between the additional bits which determines the simulation outcome. Due to space constraints we can not go into details here but would like to mention it as another empirical confirmation of the significance of the theoretical findings for the predictability of GA simulation results.

4 Discussion and conclusions

The simulations of the Cobweb model presented in this paper show that the outcomes depend crucially on the particular coding scheme. However, we also showed that the model builder is not restricted to observing these different outcomes but, by using the mathematical analysis of genetic learning in economic systems, may predict the implications of the actual simulation setup he is using. We illustrated how the long run behavior of

[15] Note that the population state corresponding to the heterogeneous equilibrium is "almost" uniform since the strings in the population vary only in one of the 11 bits. Thus, the population stays near a uniform state, as predicted in Dawid and Hornik (1996).

the simulation can be changed *systematically* by a variation of the coding scheme used. Interpreting the splitting of the output decision in purely mathematical terms we may say that we have sufficiently decreased the left hand side of the inequality (10) for $j = 0$ such that this inequality is no longer satisfied. Doing this we have made a former stable state unstable with respect to the dynamics of the GA. The question arises as to whether the stability properties of any stable state can be reversed by a similar procedure; this is answered in Dawid (1996b) where it is shown that the strict economic equilibria are the only stable uniform states the stability properties of which do not depend on the coding scheme used. Further we would like to point out that proposition 1 implies that the stability properties of non-equilibrium states may also be systematically changed by rescaling the fitness values[16] or choosing different crossover operators.

In our model we basically have a one-to-one correspondence between phenotype and genotype, i.e. differing strings induce differing actions (the only exception occurs in the model with exit and entry decisions, where two strings with entry 0 in the first bit encode the same action, even if their other bits differ in value). This does not necessarily hold in other models. Recently GAs with 'genetic waste', (i.e. some bits in the string have no influence on the phenotype) were introduced (Novkovic and Sverko, 1997). Here proposition 1 implies that we cannot expect convergence on a genotype level. However, considering the induced dynamics where the state space includes only the payoff relevant bits, we can reestablish the result of proposition 1 on a phenotype level. Differences occur in the stability criterion (10), where the left hand side has to be adapted because the probability that the crossover point lies between two relevant bit positions depends on the number of 'waste bits' in-between and is no longer uniform. In particular this implies that a strict economic equilibrium is locally asymptotically stable on a phenotype level, even if the genotype representation is not unique and, accordingly, there exists no asymptotically stable uniform population state in a genotypic sense.

Taking this into account, we have to be aware of the fact that simulation results may crucially depend on implementation details which have hardly any economic meaning. Using our stability criterion, it is at least possible to specify exactly which kind of change in the simulation setup will change the basic results of a simulation and which will not. In our opinion these insights are of great help in designing appropriate GA setups and, perhaps even more important, in the assessment of the economic implications of the simulation results.

In the GA's-as-optimization literature the relationship between coding of strategies and convergence properties seems to be well understood. However, to our knowledge this is the first paper where this point is systematically analyzed with special emphasis on economic systems. Benaroch (1996) states that, "*[S]o far, almost without exception, the very few cases*

[16] This comes as no surprise since it is well known that rescaling of the fitness may influence the stability properties with respect to the discrete time replicator dynamics (see e.g. Weissing, 1991).

where AI methods were used quite successfully in economics are the ones where economists took the lead and learned how to apply these methods and adapt them to specific economic problems." (p. 602). We believe that the insights presented in this paper may help to tailor a Genetic Algorithm to a specific economic problem.

Appendix

The mathematical model of the GA

In this section we give an exact mathematical description of the behavior of GAs in systems with a state dependent fitness function. An exhaustive derivation and analysis of this mathematical theory is provided in Dawid (1994) and Dawid and Hornik (1996).

Let again P_t denote the population at time t consisting of n binary strings of length l and the Ω the set of all binary strings of length l. We identify each element k of Ω with the integer $\sum_{i=1}^{l} k(i)2^{i-1}$, where $k(i)$ is the value of the i-th bit of k. We write S for the set of all possible population states, and Φ_t for the state of the population at time t (which a priori of course is a random vector).

For mathematical reasons we now assume that the fitness function is defined on the whole simplex Δ^r and is positive, continuous and continuously differentiable.

Considering the effect of the three operators described in Section 2 it is quite obvious that the distribution of Φ_{t+1} given $\{\Phi_0, \ldots, \Phi_t\}$ depends only on Φ_t. Therefore, we conclude that the stochastic process $\{\Phi_t\}_{t=0}^{\infty}$ is a Markov process. We denote the transition matrix of this process by $Q = [q_{\phi\phi'}]_{\phi,\phi' \in S}$. Due to the fact that we assume that both offspring from crossover are inserted into the next population, the calculation of $q_{\phi\phi'}$ is quite cumbersome and will be omitted here. Nevertheless, it is possible to derive results concerning the long run behavior of the process. In Dawid and Hornik (1996) it is shown that there exists a unique stationary distribution of the process and that this stationary distribution is concentrated around the uniform states for small mutation probabilities. However, the analysis of the Markov chain provides no information as to which uniform states are reached with high probability and which ones only with small probability.

To answer this question we study a deterministic dynamical system which provides a good approximation of the actual behavior of the stochastic process for large populations. Using a kind of "mean field theory" approach we use a system of the form

$$\phi_{t+1} = \mathcal{G}(\phi_t) \ , \tag{11}$$

where $\mathcal{G}$ is defined by

$$\mathcal{G}(\phi) = \mathbb{E}(\Phi_{t+1} | \Phi_t = \phi) \ .$$

The state space of the dynamical system is the limit of S for $n \to \infty$, namely the whole r-dimensional simplex Δ^r.

Studying this system we decompose $\mathscr{G}(\cdot)$ in two operators, the selection operator $\mathscr{S}(\cdot)$ and the mixing operator $\mathscr{M}(\cdot)$. It can be easily seen that if the state of the population at time t is $\phi \in \Delta^r$ the expected state of the mating pool is given by

$$\mathscr{S}(\phi) = \frac{\mathrm{diag}(f(\phi))\phi}{f(\phi)'\phi} \quad , \tag{12}$$

where $'$ denotes transpose. The selection operator describing the effect of proportional selection has the same form as the well known replicator dynamics in discrete time (see Weissing, 1991). This is of course no surprise as the replicator dynamics describes the evolution of a large population with the underlying assumption that the number of offsprings is proportional to the fitness of an individual.

The effect of crossover and mutation is represented by the mixing operator $\mathscr{M} : \Delta^r \mapsto \Delta^r$. The exact term for this non-linear operator is not given here but may be obtained from Dawid (1994) or Dawid and Hornik (1996).

Using this notation, we finally have

$$\mathscr{G}(\phi) = \mathscr{M}(\mathscr{S}(\phi)) \ . \tag{13}$$

Of course, the operator $\mathscr{G}$ depends on the parameters χ and μ of the GA. It is also shown in Dawid and Hornik (1996) that for $n \to \infty$ and $\mu \to 0$ the trajectory of the corresponding Markov process converges in probability to the trajectory of this dynamical system with $\mu = 0$ on every finite time interval. In other words, if the population is large and the mutation probability is small this dynamical system is with high probability within some small distance of the actual state of the population. Putting together these results with the considerations in Section 2 we may say that (13) describes the learning behavior of a population of boundedly rational economic agents.

References

Alchian AA (1950) Uncertainty, evolution, and economic theory. Journal of Political Economy 58: 211–222

Andreoni J, Miller JH (1995) Auctions with artificial adaptive agents. Games and Economic Behavior 10: 39–64

Arifovic J (1994) Genetic algorithm learning and the cobweb model. Journal of Economic Dynamics and Control 18: 3–28

Arifovic J (1995) Genetic algorithms and inflationary economies. Journal of Monetary Economics 36: 219–243

Arifovic J (1996) The behavior of the exchange rate in the genetic algorithm and experimental economies. Journal of Political Economy 104: 510–541

Arifovic J, Eaton C (1995) Coordination via genetic learning. Computational Economics 8: 181–203

Arthur WB (1991) Designing economic agents that act like human agents. A behavioral approach to bounded rationality. American Economic Review 81: 353–359

Arthur WB (1993) On designing economic agents that behave like human agents. Journal of Evolutionary Economics 3: 1–22

Axelrod R (1987) The evolution of strategies in the iterated Prisoner's dilemma. In: Davis LD (ed) Genetic algorithms and simulated annealing, research notes in artificial intelligence. Morgan Kaufmann, Los Altos
Benaroch M (1996) Artificial intelligence in economics. Truth and dare. Journal of Economic Dynamics and Control 20: 601–605
Bray M (1982) Learning, estimation and the stability of rational expectations. Journal of Economic Theory 26: 318–339
Bray MM, Savin NE (1986) Rational expectations equilibria, learning, and model specification. Econometrica 54: 1129–1160
Bray M, Kreps DM (1987) Rational learning and rational expectations. In: Feiwel GR (ed) Arrow and the ascent of modern economic theory. New York University Press, New York
Carlson JA (1968) An invariably stable cobweb model. Review of Economic Studies 2: 360–362
Dawid H (1994) A markov chain analysis of genetic algorithms with a state dependent fitness function. Complex Systems 8: 407–417
Dawid H (1996a) Learning of cycles and sunspot equilibria by genetic algorithms. Journal of Evolutionary Economics 6: 361–373
Dawid H (1996b) Adaptive learning by genetic algorithms: Analytical results and applications to economic models. Lecture Notes in Economics and Mathematical Systems, No. 441. Springer, Berlin Heidelberg New York
Dawid H, Hornik K (1995) The dynamics of genetic algorithms in interactive environments. Journal of Network and Computer Applications 19: 5–19
Dawid H, Mehlmann A (1996) Genetic learning in strategic form games. Complexity 1(5): 51–59
Day RH (1994) Complex economic dynamics, Volume I. An introduction to dynamical systems and market mechanisms, MIT Press, Cambridge, MA
Dosi G, Nelson RR (1994) An introduction to evolutionary theories in economics. Journal of Evolutionary Economics 4: 153–172
Ekeziel M (1938) The cobweb theorem. Quarterly Journal of Economics 52: 255–280
Friedman D (1991) Evolutionary games in economics. Econometrica 59: 637–666
Goldberg DE (1989) Genetic algorithms in search, optimization and machine learning. Addison-Wesley, Reading, MA
Goldberg DE (1995) The existential pleasure of genetic algorithms. In: Winter, Perioux J, Galan M, Cuesta P (eds) Genetic algorithms in engineering and computer science, pp 23–31. Wiley, Chichester
Hallagan W, Joerding W (1983), Polymorphic equilibrium in advertising. Bell Journal of Economics 14: 191–201
Ho TH (1996) Finite automata play repeated prisoner's dilemma with information processing costs. Journal of Economic Dynamics and Control 20: 173–207
Holland JH (1975) Adaptation in natural and artificial. The University of Michigan Press, Ann Arbor, MI
Holland JH, Miller JH (1991) Artificial adaptive agents in economic theory. American Economic Review 81: 365–370
Kirman A, Salmon M (1995) Learning and rationality in economics. Blackwell, Oxford
Kopel M (1997) Improving the performance of an economic system: controlling chaos. Journal of Evolutionary Economics 7: 269–289
Leontief WW (1934) Verzögerte Angebotsanpassung und partielles Gleichgewicht. Zeitschrift für Nationalökonomie 5: 670–676
Marimon R, McGrattan E, Sargent TJ (1990) Money as a medium of exchange in an economy with artificially intelligent agents. Journal of Economic Dynamics and Control 14: 329–373
Marimon R, Sunder S (1994) Expectations and learning under alternative monetary regimes: an experimental approach. Economic Theory 4: 131–162

Marks RE (1992) Breeding hybrid strategies. Optimal behavior for oligopolists. Journal of Evolutionary Economics 2: 17–38
Midgley DF, Marks RE, Cooper LG (1997) Breeding competitive strategies. Management Science 43: 257–275
Miller JH (1996) The coevolution of automata in the repeated prisoner's dilemma. Journal of Economic Behavior and Organization 29: 87–112
Mitchell M (1995) An introduction to genetic algorithms. MIT Press, Cambridge, MA
Mühlenbein H (1991) Darwin's contingent cycle theory and its simulation by the prisoner's dilemma. Complex Systems 5: 459–478
Muth JF (1961) Rational expectations and the theory of price movements. Econometrica 29: 315–335
Nerlove M (1958) Adaptive expectations and cobweb phenomena. Quarterly Journal of Economics 72: 227–240
Novkovic S, Sverko D (1997) Modeling heterogeneous agents with induced population diversity. Proceedings of the Workshop on Economics with Heterogeneous Interacting Agents, Ancona
Nyarko Y, Woodford M, Yannelis NC (1994) Bounded rationality and learning. Special Issue of Economic Theory 4: 1994
Sargent TJ (1993) Bounded rationality in macroeconomics. Clarendon Press, Oxford
Schlag KH (1998) Why imitate and if so, how? A bounded rational approach to multi-armed bandits. Journal of Economic Theory (forthcoming)
Selten R (1975) Reexamination of the perfectness concept for equilibrium points in extensive games. International Journal of Game Theory 4: 25–55
Vega-Redondo F (1995) Expectation, drift and volatility in evolutionary games. Games and Economic Behavior 11: 391–412
Weissing FJ (1991) Evolutionary stability and dynamic stability in a class of evolutionary normal form games. In: Selten R (ed) Game equilibrium model I. Springer, Berlin Heidelberg New York