

Genetic Learning as an Explanation of Stylized Facts of Foreign Exchange Markets

Thomas Lux^{1,*} and Sascha Schornstein²

¹ Department of Economics
University of Kiel
Olshausenstr. 40
24118 Kiel, Germany,

² London School of Economics
and Political Science
Department of Economics
Houghton Street
UK London WC2A 2AE

Abstract: This paper revisits the Kareken-Wallace model of exchange rate formation in a two-country overlapping generations world. Following the seminal paper by Arifovic (*Journal of Political Economy*, 104, 1996, 510 – 541) we investigate a dynamic version of the model in which agents' decision rules are updated using genetic algorithms. Our main interest is in whether the equilibrium dynamics resulting from this learning process helps to explain the main stylized facts of free-floating exchange rates (unit roots in levels together with fat tails in returns and volatility clustering). Our time series analysis of simulated data indicates that for particular parameterizations, the characteristics of the exchange rate dynamics are, in fact, very similar to those of empirical data. The similarity appears to be quite insensitive with respect to some of the ingredients of the GA algorithm (i.e. utility-based versus rank-based or tournament selection, binary or real coding). However, appearance or not of realistic time series characteristics depends crucially on the mutation probability (which should be low) and the number of agents (not more than about 1000). With a larger population, this collective learning dynamics loses its realistic appearance and instead exhibits regular periodic oscillations of the agents' choice variables.

Keywords: learning, genetic algorithms, exchange rate dynamics

JEL classification: D83, D 84, F31

* The authors are grateful for many helpful comments by the participants at the following events: the staff seminars of the Department of Economics at the Universities of Hannover and Kiel, workshops and conferences at Trieste (WEHIA '02), Swiss Exchange Zurich (Workshop on Evolutionary Finance), Leiden (Workshop on Economic Dynamics), Aix-en-Provence (SCE Annual Conference '02), and the Bundesbank Training Centre at Eltville. T.L. also wishes to acknowledge financial support by the Landeszentralbank Schleswig-Holstein.

Corresponding author: Thomas Lux, e-mail: lux@bwl.uni-kiel.de

1. Introduction

Foreign exchange markets as well as other financial markets are characterized by a number of striking ubiquitous time series features. Most prominently, (log) exchange rates seem to be non-stationary while their first differences are stationary. More precisely, unit-root tests are typically unable to reject the null hypothesis of a first-order autoregressive process with a coefficient equal to unity. This finding squares with the well-known result of Meese and Rogoff (1983) that random walk forecasts produce a lower mean-squared error in out-of-sample prediction than reduced-form structural models of macroeconomic fundamentals. It has been argued that these findings can be explained by *speculative efficiency* of foreign exchange markets, which simply means one interprets the foreign exchange market as an informationally efficient market in the sense of the Efficient Market Hypothesis.

While from this perspective the unit-root property may not be viewed as a conundrum, other well-known features have defied straightforward explanations until recently. The most pervasive ones are the fat-tail property of relative price changes and the clustering of volatility in these time series. Traces of these features are easily recognizable in all records of high-frequency data (probably up to weekly frequency) of foreign exchange markets (to our knowledge, without any known exception). The fat-tail property implies that the unconditional distribution of daily returns (as well as those of higher and somewhat lower frequency) has more probability mass in the tails and the center than the standard Normal distribution. This also means that extreme changes occur more often than would be expected under the assumption of Normality of relative daily price changes. Volatility clustering means that periods of quiescence and turbulence tend to cluster together. Hence, the volatility (conditional variance) of exchange rate changes is not homogeneous over time, but is itself subject to temporal variation.

Explanations of these stylized facts have been elusive until very recently. Perhaps, the silence of economic theory on this issue is not too surprising given that the above regularities are features of time series as *a whole* and, hence, could only be explained by dynamic models of the evolution of the trading process in the pertinent market. From the viewpoint of informational efficiency, the characteristics of returns would, of course, have to be explained by similar characteristics of the *news arrival process*, but due to the unobservability of the later, this hypothesis can hardly be subjected to econometric scrutiny. As an alternative, some authors have recently argued that fat tails and clustered volatility can be obtained as a result of interactions of heterogeneous economic agents. Examples of this emergent literature include Lux and Marchesi (1999, 2000), Chen, Lux and Marchesi (2001), Kirman and Teysiere (2001), Gaunersdorfer and Hommes (2000), Chiarella and He (2001), Iori (2002) and Bornholdt (2001). Lux and Marchesi, Gaunersdorfer and Hommes, and Chiarella and He have models of fundamentalist - chartist interaction in financial markets which give

rise to realistic behavior of the resulting time series (in terms of the above stylized facts). In Lux and Marchesi and Gaunersdorfer and Hommes, the authors try to provide some hints of general mechanisms that could generate these time series properties irrespective of the details of their exemplary models. In the former case, it is a critical behavior of the dynamics in the vicinity of a continuum of equilibria with an indeterminate composition of the population in terms of strategies pursued by individuals. Gaunersdorfer and Hommes get similar dynamics from a model with co-existing attractors in which noise leads to switches between different states. Still different mechanisms prevail in Iori (2002) and Bornholdt (2001) who use lattice-based structures for modeling the interactions among traders. Interestingly, a recent paper by Arifovic and Gencay (2000) on an artificial currency market with genetic learning of strategies also suggests emergence of realistic features of the resulting exchange rate dynamics (cf. Fig. 1). However, they do not provide a detailed analysis concerning the above properties. One of the aims of this paper is to fill this gap. In particular, we will try to quantitatively assess the *degree* of fat tailedness and volatility clustering this model generates. We are also interested in the sensitivity of these quantitative measures with respect to key parameters of the model. To get an impression of the sensitivity with respect to parameter variations, we will try to figure out how the time series properties depend on the genetic algorithm parameter and the number of agents populating the market (as will probably become clear in the presentation of the model, the values of the few economic variables of the model are of lesser importance in this respect). We then relate our findings to those obtained for other models of artificial financial markets and try to provide an explanation for the crucial importance of the number of individuals for the qualitative outcome of the model.

Our analysis proceeds in the following steps: sec. 2 will introduce the underlying model of the foreign exchange market, the well-known Kareken-Wallace two-country overlapping generations model. Sec. 3 gives details on the genetic algorithms which we apply to model the learning of our agents. In sec. 4, we review the statistics used for assessing how realistic the model's output is. Sec. 5 presents the results of extensive Monte Carlo work, and sec. 6 tries to provide an explanation for the surprising behavior we find in the case of a very large population. Sec. 7 concludes.

Fig. 1 about here

2. The Kareken and Wallace OLG Economy

As a version of the Kareken and Wallace (1981) two country model, the underlying economic structure is extremely simple: at each date t , one-half of the entire population is replaced by a new generation (the young), while the remaining members are in the second and final period of their lives (the old) and will be replaced in the next period by another young generation. Each agent is endowed with w_1 units of a homogeneous good in its first period and with w_2 units in the second period of its

live. There is neither production nor inheritance of goods. Intertemporal consumption smoothing can be achieved via money holdings of currency of the home and foreign country.

With identical preferences of all agents, $U(c_i(t), c_i(t+1))$, their consumption plans and money demand are derived from

$$(1) \quad \max U(c_i(t), c_i(t+1))$$

subject to: $c_i(t) \leq w_1 - \frac{m_{i,1}(t)}{p_1(t)} - \frac{m_{i,2}(t)}{p_2(t)}$ and

$$c_i(t+1) \leq w_2 + \frac{m_{i,1}(t)}{p_1(t+1)} + \frac{m_{i,2}(t)}{p_2(t+1)}.$$

With $c_i(t)$, $m_{i,1}(t)$ and $m_{i,2}(t)$: consumption and money holdings of agent i ($i=1, \dots, N$) at time t , w_1 and w_2 the homogeneous endowment levels, and $p_1(t)$ and $p_2(t)$ the price levels in both countries at time t . Note that with this setup, it even does not matter, how many of the agents are citizens of countries 1 or 2 as their economic decisions are not affected by their geographic location.

Assuming that nominal money supply H_1 and H_2 , is constant, and denoting by $s_i(t)$ overall ‘savings’ of individual i , the price levels at time t are determined by:

$$(2) \quad p_1(t) = \frac{H_1}{\sum_i f_i(t) \cdot s_i(t)}, \quad p_2(t) = \frac{H_2}{\sum_i (1 - f_i(t)) \cdot s_i(t)} \quad \text{with}$$

$$f_i(t) = \frac{m_{i,1}(t) / p_1(t)}{s_i(t)} : \text{the fraction of currency 1 in agent } i \text{'s portfolio at time } t.$$

The exchange rate, $e(t)$, between both countries is, then, obtained as $e(t) = \frac{p_1(t)}{p_2(t)}$.

It is straightforward to spell out some results on possible equilibria (steady states) of this model. Using $c(t)$ and $f(t)$ as choice variables of our agents, it is immediately obvious that a stationary solution requires the rates of returns on both currencies to be equal, i.e. $\frac{p_1(t)}{p_1(t+1)} = \frac{p_2(t)}{p_2(t+1)}$.

Simple inspection shows that such a steady state has the following properties:

(1) the exchange rate is constant over time: $e(t+1) = e(t) = e^*$. However, the *level* of the exchange rate in equilibrium, e^* , is indeterminate and may be any value in the half-line, $e^* \in (0, \infty)$.

(2) Given the equality of returns from both currencies, the equilibrium composition of each agent's portfolio is indeterminate as well and f_i^* , therefore, might assume any admissible value $f_i^* \in [0,1]$. Note that we do not have to assume identical portfolio choices of the agents in equilibrium. Any constant distribution of the f_i^* 's over generations would be consistent with constant rates of returns and a constant exchange rate. We could even allow for certain changes of the distribution of the f_i^* 's: equilibrium would still prevail as long as the mean value of their distributions remains the same over time.

(3) Since the rates of returns from both assets are identical in equilibrium, optimal consumption plans are independent from the portfolio composition. With a well-behaved utility function, utility maximizing choice c_i^* will typically be unique and identical for all individuals with the same utility function and endowment structure. For example, with utility functions $U_i = c_i(t) c_i(t+1)$ and constant money supply in both countries, rates of return would be equal to one and c_i^* would be $0.5(w_1 + w_2)$ for all individuals.

The non-uniqueness of the equilibrium exchange rate in this type of model derives from the absence of typical macroeconomic fundamentals of monetary models. In view of the evidence on unit-roots in empirical data, this feature has been emphasized as an advantage of their model by Kareken and Wallace. One could indeed imagine that added random fluctuations could easily produce a unit-root dynamics, since random disturbances could lead to a random motion of the exchange rate along the continuum of possible equilibria (every time, the equilibrium is distorted by random shocks, the exchange rate would settle at a new equilibrium). However, non-uniqueness of equilibria also raises the questions of selection of equilibria and coordination of agents. These questions have been taken up first by Sargent (1993) who modeled learning via stochastic approximation algorithms. Later on Arifovic (1996) considered GA learning in the Kareken–Wallace framework. Looking at the evolution of returns instead of the level of the exchange rate Arifovic and Gencay (2000) recovered realistic features in the continuing fluctuations of the resulting dynamics.

3. Genetic Algorithm Learning

Genetic algorithms have been introduced by Holland (1975) as a stochastic search algorithm for numerical optimization. This approach uses operations similar to genetic processes of biological organisms to develop better solutions of an optimization problem from an existing 'population' of randomly initiated candidate solutions. Typically, the proposed solutions have been encoded in strings (chromosomes) using a binary alphabet (see Dawid, 1999 for a general introduction). This is also the structure of the GAs applied in Arifovic and Gencay (2000). Each individual's decisions are encoded in a binary string of length $l=30$, whose first twenty elements encode first-period consumption and whose remaining ten entries encode the fraction of currency one in his portfolio.¹

¹ Choosing $l = 30$ with substrings of twenty and ten bits, respectively, we closely followed Arifovic (1996) and

With $a_{i,t}^k$ denoting the value at the k -th position of the string (0 or 1), the binary string is translated into a real-valued number in the following way:

$$(3) \quad c_i(t) = \sum_{k=1}^{20} a_{i,t}^k \cdot \frac{2^{k-1}}{K_1}, \quad f_i(t) = \sum_{k=21}^{30} a_{i,t}^k \cdot \frac{2^{k-21}}{K_2}$$

where K_1 and K_2 are normalizing constants to restrict the possible real values to a predetermined admissible interval.

In order to have $f_i(t) \in [0,1]$, K_2 is set equal to $K_2 = 2^{10} - 1$, while $c_i(t)$ should be within the interval $[0, w_1]$ to guarantee viable consumption plans. With $w_1 = 10$ in Arifovic and Gencay's simulations, this amounts to $K_1 = \frac{2^{20} - 1}{10}$.

The overlapping generations structure of the model implies an overlapping genetic algorithms structure of the evolutionary process. After each period, half of the population members have completed their life cycle. With the resulting consumption in their old age their achieved fitness (utility) can be determined and used for the genetic creation of a new pool of agents entering the economy as the young generation of the following period.

The genetic operations applied in this step are the following:

(1) **Reproduction**: from the pool of old individuals, copies are selected (with replacement) with probabilities depending on their relative fitness, i.e. on $\frac{U_i}{\sum_i U_i}$.

Other algorithms for reproduction could be chosen as well: proposals in the GA literature include rank-based reproduction in which only the rank (not the absolute fitness) determines the probability of reproduction, and tournament selection, in which one repeatedly draws n_1 (say 5) individuals from the pool and accepts the $n_2 < n_1$ (say 2) with highest utility among them for the new generation. Below we report only results for fitness-based reproduction. Experiments with rank-based and tournament selection have also been carried out yielding almost identical results.

(2) **Crossover**: when the pool of potential new members of a generation is complete, genetic material is exchanged between them. The simplest way is randomly selecting a pair of parent strings and swapping genetic material (bits) between both chromosomes. Here, we again follow the algorithm used by Arifovic and Gencay in selecting randomly an integer in the range of $[1,29]$ and constructing offspring by combining the genetic material from the left of this position from parent one

Arifovic and Gencay (2000). However, like with most other details of the genetic algorithm implementation, variation of these numbers did not change the qualitative characteristics of the dynamics.

with that from the right-hand part of parent two and *vice versa*. Note that the cross-over operation is carried out with a certain probability p_{cross} only, while with probability $1-p_{\text{cross}}$ the offspring are unchanged copies of their parents. Alternative implementations of the crossover operator include two-point crossover (exchange of material in an interval between two randomly chosen bits) and uniform cross-over (the two offsprings are random recombinations of their parents' bits). Again, the whole chain of our findings reported below seems to be robust with respect to the choice of the particular cross-over operator.

(3) **Mutation** simply means that each position within a string is altered with a certain probability p_{mut} to the other value of the binary alphabet.

(4) Finally, the **election** operator tests newly generated offspring before letting them enter the population. In order to avoid a decrease of the fitness of the overall population due to the genetic alteration of strategies, only those among the offspring are accepted which are at least as fit as one of their parents. If after crossover and mutation offspring have lower fitness, exact copies of their parents are placed into the new generation.

Beside these traditional binary coded GAs, we also experimented with *real coded* GAs in the present framework. A real coded GA simply uses a real representation of the choice variables. In our case, a real-coded chromosome would, therefore, consist of a pair $\{c_i(t), f_i(t)\}$.

Mutatis mutandis, similar genetic operations can be defined for this variant (cf. Herrera *et al.*, 1998, for an overview on real-coded GAs). First, reproduction occurs in the same way as with binary GAs. As for cross-over, a number of alternative mechanisms have been proposed in the literature. Here, we follow Eshelman and Shaffer (1993) in picking new choice variables which are uniformly and independently drawn from an interval covering the pertinent values of the parents' chromosome. To illustrate let $\bar{c} = \max\{c_A(t), c_B(t)\}$, $\underline{c} = \min\{c_A(t), c_B(t)\}$, and $d_c = \bar{c} - \underline{c}$ with $c_A(t)$ and $c_B(t)$ consumption choices of parents $i = A, B$. Then, the first-period consumption of offsprings is determined by uniform random draws from $[\underline{c} - \gamma d_c, \bar{c} + \gamma d_c]$. A similar operation yields the new portfolio fractions of the offsprings. Note that γ is a predetermined value that allows for some 'experimentation' within regions not covered by the genetic material of the parents. Its role is also to compensate for the drift towards the mean of the admissible strategy space from a crossover operator with $\gamma = 0$. Herrera *et al.* (1998) show that this algorithm has better performance on some test problems than many alternatives. Mutation with real variables is done by using Normal random variables with mean zero and small variance to slightly change the prevailing choice variables. Election, finally, occurs in the same way as with binary coding.²

² On a first view, the real coded GAs seem to be a much more natural way to deal with any real-valued problem than binary GAs. However, good reasons are given in the literature to actually prefer binary coding in many applications. In particular, although at first view GAs seem to process only the particular chromosomes within the population, they actually allow for a parallel processing of many different parts of alternative solutions. This might be sensed by considering the following example: the binary coded values for 0 and 4, i.e. 000 and 100, have two – thirds of their genetic material in common, so that it is 'easy' for the genetic operations to switch from one

4. Empirical Benchmarks

To see whether our foreign exchange market populated by genetically learning overlapping generations has realistic time series properties we use a battery of statistical tests.

(i) **Unit root tests:** a realistic market should yield a exchange rate dynamics which appears to be close to a random walk. We, therefore, perform typical tests for the presence of a unit root in our synthetic time series using the standard Dickey Fuller (DF) and Augmented Dickey Fuller (ADF) tests. The underlying data are logs of the exchange rate since from the symmetric construction of the Kareken-Wallace model we would also expect symmetry of relative changes. For the ADF test, we also included the first three differences as independent variables.

(ii) **Fat-tail property:** the recent econometric literature has provided a very sharp characterization of this feature. In particular, it could be shown that the decline of the probability mass in the tails follows a power-law with a coefficient that is astonishingly uniform across markets. This amounts to large returns (r_t) following a relationship: $\text{Pr ob.}(|r_t| > x) \sim x^{-\alpha}$ with the so-called ‘tail index’ α hovering within the interval between 2 and 5. Estimation of α using conditional maximum likelihood is straightforward, and a wealth of supporting evidence can be found in De Vries (1994) and Dacorogna *et al.* (2001). As a typical example, estimation of the tail index for the DM/U.S.\$ exchange rate with daily data ranging from 1974 to 1998 yields an estimate $\alpha = 3.69$ (95% confidence interval: 3.38 to 4.10) when using the five percent largest absolute returns. The review by De Vries and the monograph by Dacorogna *et al.* give similar statistics for other currencies.

(iii) **Volatility clustering:** this feature can be characterized by autoregressive dependence in various measures of volatility. Here we also have a very precise and uniform picture from almost all available data sets. In particular, it has been found that the dependence in volatility measures like squared or absolute returns extends over very long time horizons and exhibits a hyperbolic decay of the autocovariance function: $E[x_t x_{t-\Delta t}] \sim \Delta t^{-\kappa}$ with x_t : squared or absolute returns. This slow decay is in contrast to fast (exponential) decline and is also denoted as long-term dependence. Like with the fat-tail property, quantitative measurements of the decay parameter κ give very uniform results across markets. As a benchmark for our later analysis of simulated data, we give estimates from the frequently used periodogram regression technique due to Geweke and Porter-Hudak (1983). These authors devise a method for estimation the *parameter of fractional differentiation*, denoted d in the following, which is related to κ by: $\kappa = 1 - 2d$. An estimate of d significantly larger than 0 would show evidence for the long-memory property. Inability of rejection of $d = 0$ would indicate absence of long-term dependence. For the sake of illustration, daily DM/\$ data yield (95%

to the other. With real-coded GAs, 0 and 4 have nothing in common and a large succession of crossovers and mutations is needed to move from one to the other.

confidence intervals in brackets): raw returns: $d = 0.07$ (-0.09, 0.23); squared returns: $d = 0.24$ (0.08, 0.40), absolute returns: 0.29 (0.13, 0.45). Again, a glance at, for example, Dacorogna *et al.* shows that these figures are quite representative for foreign exchange data (as well as for financial data in general). The inability of rejection of $d = 0$ for raw returns, of course, squares well with the unit-root property of log exchange rates. The finding of a higher level of persistence in absolute returns rather than squared returns is also quite universal and has by itself motivated a large body of recent econometric literature.

5. Pseudo-Empirical Results

Tables 1 to 6 present the results of a large number of experiments with various versions of our artificial economy. Tables 1 to 3 show results for binary coded GAs, while tables 4 to 6 are concerned with simulations using real coded GAs.

Our starting point was the scenario underlying the simulations by Arifovic and Gencay (2000). The particular Kareken-Wallace economy in this paper had the following properties: all individuals share a common utility function $U_i = c_i(t) c_i(t+1)$, endowments are $w_1 = 10$, $w_2 = 4$, and nominal money supplies are $H_1 = 3000$, $H_2 = 3600$. It can immediately be seen that this leads to a steady state consumption level of $c^* = c(t) = c(t+1) = 7$ and steady-state savings $s^* = s(t) = s(t+1) = 3$. *Prima facie*, it appears unlikely that changes in these economic variables should yield greatly different results (as long as endowments would lead to positive savings). Experimentation with different parameters and alternative utility functions (e.g., logarithmic utility) confirmed this conjecture.

In fact, our interest here is more (i) in the sensitivity with respect to the details of the learning dynamics, and (ii) the influence of the number of agents. Our interest in the effects of the size of the market derives from some puzzling earlier findings. Namely, a number of studies have revealed that existing multi-agent models of financial markets lose their realistic time series properties when increasing the number of agents (Egenter *et al.*, 1999; Yeh, 2001, Challet and Marsili, 2002). Since published work on artificial markets with GA learning has used only a very limited number of agents, typically below 100, it seems worthwhile to explore the behavior of larger economies.

Let us start with the effects of varying the GA's parameter settings. From the two parameters of the binary genetic operations, p_{mut} and p_{cross} , we found the first to be the more interesting one in that variation of p_{cross} only led to slight variations of the statistical properties. In a first set of experiments we, therefore, fixed p_{cross} at 0.6 and also fixed the population size at the level used in Arifovic and Gencay, $N = 60$ (i.e., 30 individuals in each generation). In order to see the effects of variation of p_{mut} , we varied this parameter from 0.005 to 0.05 (with increments of 0.005) and applied the statistical analyses outlined above to 100 samples each containing 2,000 data points (a length of the

data series comparable to many empirical records). The 100 samples are taken as non-overlapping windows from a simulation continuing over 205,000 periods where the first 5,000 data points have been discarded to account for transient behavior. Minimum, median and maximum of the tail index estimates (for tail sizes of 2,5%, 5% and 10% of the data) are shown in Table 1, while the minimum, median and maximum of the estimates of d are given in the Table 2. Table 3, finally gives the median and range of estimates of the autoregressive parameter from the Dickey-Fuller and Augmented Dickey-Fuller test together with the frequencies of one-sided (in brackets: two-sided) rejection of the unit-root null under a 95% confidence level. Results are quite homogeneous with respect to our three stylized facts in so far as the behavior is most realistic for small values of p_{mut} around 0.05 to 0.01. In this region we have a high percentage of non-rejection of a unit root in log exchange rates (at least for the ADF test which corrects for short-run dynamics) together with median values of d close to their empirical counterparts. The median tail index estimates might appear somewhat too small, but are still within the range observed with empirical data. Fig. 1 illustrates that returns obtained with this setting of the GA parameters indeed do look very realistic and may be hard to distinguish from real-life records with the naked eye (at least, after, proper adjustment of the scale of the fluctuations). However, when increasing p_{mut} beyond 0.02, rejection of a unit root in favour of a root smaller than unity occurs in all cases, the fractional differencing parameter for raw returns becomes negative (which is also a signature of mean reversion), and the temporal dependence in squared and absolute returns declines. Finally, the tail index becomes somewhat too high.

Tables 1 to 3 about here

In our second set of experiments, we then varied N keeping the mutation probability at the value 0.01. Since the simulations become more time-consuming with increasing N , we restricted our investigation to 25 samples (i.e. a time series of overall length of 55,000 time steps for each parameter set). Here the changes are even more dramatic. When moving from small ($N = 20$) to very large markets ($N = 10,000$), we get an even larger drop of the autoregressive coefficient in the unit-root tests, highly negative d 's for raw returns and a total fading out of volatility clustering (the d 's of squared and absolute returns approaching zero). The tail index decreases and has median values below 2 for the maximum size of the market ($N = 10,000$).

Before turning to explanations, let us look at the pertinent results for real-coded GA's for which we also varied both the set-up of the mutation operator and the number of agents (Tables 4 to 6). Again, the parameters of the crossover operator are kept constant (uniform crossover as described in s. 3 with a parameter $\gamma = 0.2$ was used). The mutation operator now has two parameters, the probability of its activation p_{mut} and the variance of the Normal mutations, σ_{mut} .

The upper and middle part of Tables 4 to 6 exhibit the effects of systematic variation of p_{mut} and σ_{mut} . Again, 100 samples of 2,000 data points each have been used. Since markets with real-coded

GA's need less computation time, we were able to use 100 subsamples when assessing the effects of market size and could also use a somewhat larger maximum size of $N = 20,000$.

In broad harmony with the binary-coded case, realistic properties are obtained with small mutation probability and small variance of mutation. Nevertheless, certain differences are observed between the binary and real-coded cases: in particular, the DF and ADF tests are still unable to reject the unit root hypothesis in the majority of cases even with a relatively large mutation rate and a large number of traders. On the other hand, mean reverting tendencies are clearly observable in the estimates of the fractional differencing parameter for raw returns in these cases. It might be that the evolution of the system is simply very slow so that in many periods the slight variations in the exchange rate go through as a unit root process. In contrast to the binary case, in the real-coded one the two-sided test often yields more rejections than the one-sided test. The additional right-hand rejections may be related to the sudden bursts of activity visible in the upper part of Fig. 3. Overall, the much higher rate of rejection of the unit root hypothesis in the binary case might be due to the higher degree of stochasticity inherited from thirty instances of mutation (for every bit) instead of two instances only in the real-coded case. Some reflection, in fact, reveals that the values of the mutation probabilities can not be directly compared between the binary coded and real coded GAs. For example, a $p_{\text{mut}} = 0.033$ (per bit) for binary coded GAs implies that with chromosomes of thirty bits, almost every individual will undergo some mutation of its genetic material. This amounts to a much higher mutation rate *within the population* with $p_{\text{mut}} = 0.01$ in the binary case than with 0.05 for real coded bits. Except for non-rejection of the unit root hypothesis, the pattern of results is, in fact, almost the same as with binary GAs when increasing N . Looking at the resulting time series, we find in both the real-coded and binary coded case a tendency towards persistent and very regular cycles when increasing the number of agents (cf. Figs. 3 and 4).

Tables 4 to 6 about here

Besides this similarity in the results from both GA variants (and a number of alternative implementations of various operators) our experiments also show that we do not need all of the typical elements of the GA to arrive at these results. Essentially, the structure of the artificial economy remains unchanged if we dispense with both the selection and crossover operators. Mutation and election alone are capable of producing these patterns, but they are also crucial for their emergence (although selection and crossover tend to speed up convergence to regular oscillations).

6. Interpretation of Results

What is the reason for this fading out of realistic time series properties with increasing mutation

probability and increasing number of agents? It is probably not too difficult to answer the first part of the question: A high mutation probability introduces a certain tendency of mean-revision of the choice parameters which is reflected in similar mean reversion of the exchange rate. For the sake of illustration, imagine a model with a mutation probability equal to 1 in the case of binary coded GAs (corresponding to a $p_{\text{mut}} = 1$ together with a high variance of mutation in the real coded case). This would lead to a stationary random distribution of agent's strategies. All deviations from the average would be corrected by the new random choice of the population in the next period and, hence, one gets a tendency of return to the mean values of the distribution of $c_i(t)$ and $f_i(t)$. Since these choice variables determine prices and the exchange rate, mean-reversion would also carry over to these variables as well. The higher the mutation rate, the higher the influence of this tendency. Higher p_{mut} , therefore, leads to less persistence in exchange rates so that the exchange rate dynamics becomes stationary and unit roots can be rejected. This suspected change in the appearance of the time series is already well recognizable when comparing binary coded GAs with $p_{\text{mut}} = 0.01$ and 0.05 , respectively (cf. Fig. 2). While the former series (or parts of it) might be accepted as a random walk by the DF and ADF tests, the second surely will not. As a conclusion, we infer from these considerations that random experimentation with strategies has to be limited in order to get realistic appearance of the time series.

Fig. 2 about here

As to the number of agents, a glance at the time series resulting with different sizes of the market reveals some perplexing patterns (Figs. 3 and 4). What looks extremely unsystematic with a small market ($N = 200$), becomes much more regular when the market becomes larger and eventually evolves into an almost perfectly regular oscillatory motion of some 'macroscopic' variables (in particular, the exchange rate and the average fraction of domestic currency in the portfolio).³ On a close inspection, the short spikes in the exchange rate fluctuations in Fig. 3 are very similar to one half-cycle in Fig. 4. Such a regular pattern is puzzling at first view as it is the result of the evolving decisions of a very large ensemble of autonomous artificial agents. Keeping in mind that the quantity displayed in the middle part of Figs. 3 and 4 is a *population average* what these oscillations show is a systematic shift of the whole distribution of this variable within an heterogeneous ensemble of agents. To our knowledge nothing of a similar type of self-organizing patterns is known in multi-agent systems with GA learning in economics or other fields. A certain clue to the underlying mechanisms can be obtained through analysis of what happens in the case of a large economy (i.e., with the number of agents going to infinity). Noting that GAs are an adaptive adjustment scheme that drives the actual average behavior of the population towards the momentary optimum of the choice variables, the large economy case might be described via the resulting deterministic mean value

³ We have chosen a higher mutation probability compared to our benchmark case in Fig. 4 since it both leads to a decrease of the amplitude of the oscillations and provides faster convergence to almost perfectly regular patterns. However, the trend to emerging regular oscillations is also clearly visible in other simulations with either real-coded or binary GAs.

dynamics of the choice variables. Unfortunately, the present dynamics is too complicated to derive explicit dynamics laws for the large economy limit.⁴ However, some heuristic considerations will reveal most of the important elements of our dynamics.

Fig. 3 about here

As an adaptive adjustment scheme, the genetic algorithm has at its intrinsic benchmark those values $c_i^*(t)$ and $f_i^*(t)$ which would have been optimal choices for the population at time t which inherits its genetic material to the generation born at time $t+2$.

It is easy to see that for our utility function $U_i = c_i(t) c_i(t+1)$ optimal behavior of individual i at time t would have been:

$$(4.a) \quad c_i^*(t) = \frac{1}{2} \left(w_1 + \frac{w_2}{f_i(t) \frac{p_1(t)}{p_1(t+1)} + (1-f_i(t)) \frac{p_2(t)}{p_2(t+1)}} \right),$$

$$(4.b) \quad f_i^*(t) = \begin{cases} 1 \\ \text{in det er min ate} \\ 0 \end{cases} \quad \text{if} \quad \frac{p_1(t)}{p_1(t+1)} \begin{matrix} > \\ = \\ < \end{matrix} \frac{p_2(t)}{p_2(t+1)}.$$

Via prices at period $t+1$, the optimal behavior of generation t also depends on the decisions of the next generation. The election operator, in fact, guarantees that the resulting new individuals accepted after selection, crossover and mutation are at least as good as their parents. For the portfolio component, this clearly implies $f_i(t+2) \in [f_i(t), f_i^*(t)]$ while consumption might also overshoot its target, $c_i^*(t)$ as long as the resulting new parameter set provides at least the utility level enjoyed by the parent individuals. Since in any out-of-equilibrium situation, the goal value for the fraction of domestic assets will be the same, 0 or 1, for all members of the population, in the large economy limit, the motion of the mean value $\overline{f}(t)$ will follow a deterministic path towards these values as long as the pay-off differential from holding domestic or foreign currency does not change.

Fig. 4 about here

To see the source of the regular fluctuations, consider a stationary situation with homogenous choice variables $c_i(t) = c_i(t-1) = c^*$ and $f_i(t) = f_i(t-1) = f_0$, where f_0 might be any admissible value between 0 and 1. Arifovic (1996) has already emphasized that any such equilibrium of this GA economy is *evolutionary unstable* since any local disturbance (mutation) will be magnified by the

⁴ Available analytical approaches to genetic algorithm dynamics consider simpler examples and are not applicable to the present model (e.g., Prügel-Benett, 1994, or Srinivar and Patnaik, 1996).

ensuing adjustments of the remaining members of the population. For the sake of the argument, assume that only one individual undergoes a mutation when the generation $t+1$ emerges from the genetic operations on generation $t-1$. Assume furthermore that this mutation amounts to an increase of the fraction of domestic money in the portfolio of this agent, while its consumption remains unchanged at the initial equilibrium level. Since this new strategy will have the same utility as its parents (because returns from holding either currency are initially identical), the election operator will allow this offspring to replace one of its parents. However, the presence of this mutant suffices to change the structure of returns for agents of generation t : instead of equal returns, they experience a higher pay-off from holding domestic money:

$$(5) \quad R_1(t) = \frac{\sum_i f_i(t+1)(w_1 - c^*)}{\sum_i f_i(t)(w_1 - c^*)} > R_2(t) = \frac{\sum_i (1 - f_i(t+1))(w_1 - c^*)}{\sum_i (1 - f_i(t))(w_1 - c^*)}.$$

Although the differential might be very small, it suffices to make $f^* = 1$ the dominant strategy. Hence, the *random* change from generation $t-1$ to generation $t+1$ induces a *systematic* shift into the same direction when generation t inherits its genetic material to generation $t+2$. Does it also lead to changes in the consumption behavior of generation $t+2$? Changes in the momentary optimal consumption level occur if the denominator in eq. (4.a) deviates from one. With consumption still equal to its steady state level at generations t and $t+1$, this denominator amounts to $f_i(t) \frac{\overline{f(t+1)}}{\overline{f(t)}} + (1 - f_i(t)) \frac{1 - \overline{f(t+1)}}{1 - \overline{f(t)}}$. Given our assumption, $\overline{f(t+1)} > \overline{f(t)}$, it is easy to see that the denominator is $> (<) 1$ for individuals with $\overline{f(t)} > (<) \overline{f(t)}$. It would, therefore, be optimal for the former to reduce first-period consumption, while the later would find it advantageous to increase it. However, since we have also assumed that all individuals share the same choice variable $f_i(t) = f_0 = \overline{f(t)}$ at time t , the former consumption level c^* would still be optimal for all generation t members as long as their portfolio choice remains unaffected by the genetic operations, so that isolated changes of $c_i(t)$ would not survive the election operator test. Note that these considerations apply only in the case of isolated genetic changes of either $\overline{f(t)}$ or $c_i(t)$. Often both variables will be affected by the genetic processes. To see more generally, what kind of arbitrary combined genetic changes would survive in our scenario, we can take stock of the traditional concept of indifference curves. Accepting only offspring who are at least as $\overline{f(t)}$ as their parents, the election operator only allows those to enter the population whose choice variables positions them on the same indifference curve like their parents or a higher indifference curve. Consider the utility obtained by parent individual i :

$$(6) \quad U_i(t) = c_{i,1}(t) \cdot (w_2 + (w_1 - c_{i,1}(t)) \cdot \rho_i(t)), \text{ with } \rho_i(t) = f_i(t) \cdot R_1(t) + (1 - f_i(t)) \cdot R_2(t).$$

The slope of an indifference curve is given by:

$$(7) \quad \left. \frac{dc_{i,1}(t)}{df_i(t)} \right|_{U=\text{const}} = \frac{-c_{i,1}(t) \cdot (w_1 - c_{i,1}(t)) \cdot (R_1(t) - R_2(t))}{w_2 + (w_1 - 2c_{i,1}(t)) \cdot \rho_i(t)}.$$

For $R_1(t) > R_2(t)$ this gives:

$$(8) \quad \left. \frac{dc_{i,1}(t)}{df_i(t)} \right|_{U=\text{const}} \begin{matrix} > \\ < \end{matrix} 0 \quad \text{if} \quad c_{i,1}(t) \begin{matrix} > \\ < \end{matrix} c^*(f_i(t)) = 0.5(w_1 + w_2 / \rho_i(t)).$$

This yields the parabolic shape of the indifference curves exhibited in Fig. 5. In the case of $R_1(t) < R_2(t)$, the inequalities in eq. (8) are reversed and the indifference curve parabolas have the opposite orientation. In the present case, $R_1(t) > R_2(t)$, it can be inferred from eq. (7), that higher utility can only be achieved if a higher fraction of domestic assets is chosen. If, however, this necessary condition for an improvement is met, a certain range of higher or lower consumption levels would be accepted by the election operator. As can be seen from Fig. 5, in the above situation in which $\overline{f(t+1)}$ slightly exceeds $\overline{f(t)}$, the spectrum of utility improving changes is slightly asymmetric with respect to consumption. Assuming that all initial consumption levels are close to c^* , the pay-offs can be reduced to: $R_1(t) = \frac{\overline{f(t+1)}}{\overline{f(t)}}$, $R_2(t) = \frac{1 - \overline{f(t+1)}}{1 - \overline{f(t)}}$. According to the arguments given above,

individuals with a below (above) average fraction of domestic money would, then, have $c^*(f_i(t)) > (<) c^* = 0.5(w_1 + w_2)$. The expected direction of combined changes of c_i and f_i would, therefore, depend on the individual's position within the distribution of the f_i 's. For roughly half of the population on average somewhat higher consumption levels would pass the election operator, while for the other half of its members, the genetic operations would slightly favor a reduction in consumption. With a symmetric distribution of the f_i 's, the expected macroscopic effect of induced changes of c_i would be close to zero. With an asymmetric distribution, skewness would somewhat favor one or the other direction of changes, but since this is a third-order effect, one might expect it to be negligible. This conjecture is supported by our simulations which show no clear trend in the development of $c_i(t)$ over time. What can be observed, however, is that an increase of the number of agents leads to a reduction of the size of fluctuations of $c_i(t)$, cf. Figs. 3 and 4.

With only small and rather unsystematic changes of $c_i(t)$, the systematic changes of the portfolio composition will dominate the dynamics. The attraction towards the extreme solutions will, then, be self-reinforcing leading to an ever increasing fraction of domestic assets in the upward part of the cycle. Since every new round of genetic breeding of a new cohort starts at a higher average level of f_i compared to the previous period, the deterministic limit of the stochastic dynamics will also lead to a higher new average value two periods later compared to the period before. In the infinite population limit, this trend will continue until the entire population will have converged to $f_i(t) = 1$. Although in this situation, all inherent tendencies of genetic changes come to a halt, the first mutation of an individual leading to an $f_i < 1$ will destabilize this stationary state again and generate a systematic

downward trend which over time leads to a convergence of the whole population to $f_i(t) = 0$. Here again, any mutation will exert a destabilizing tendency commencing the upward part of the cycle... and so on *ad infinitum*. Note that this endlessly repeating cycle should also somehow exist as a tendency in the finite population case (since what we observe in the large population case should correspond to the pure mean value dynamics). Of course, the cyclical development shown in Fig. 4 would not be consistent with a unit root in log exchange rates (it *is* a clearly mean reverting process) and volatility clustering. In an sense, with a large population, the inherent randomness of the artificial economy gets lost and the measurable macroeconomic observables ($p_i(t)$, $e(t)$) become deterministic quantities.

However, at least for very small populations, this inherent structure of the combined genetic and economic process seems to be entirely concealed by the random elements in the genetic processes on the level of the individual. In fact, the systematic tendencies worked out above will be subject to more random distortions with a small population size. An upward or downward tendency will be inverted as soon as the portfolio fraction of a new generation is not higher (lower) than that of the preceding one. It is the more likely that this random event happens the smaller the size of the population is. The apparently realistic time series characteristics result from situations where this happens with a very high probability within a few time steps. This explains why these more irregular dynamics with recurrent bursts of activity are only observed within a certain range of small numbers of agents.

We end our attempts at providing intuitive explanations of the evolutionary dynamics of our model with some remarks on the dynamics of second moments. As can be seen from Fig. 4, even the standard deviations of our choice variables exhibit predictable systematic patterns over the cycle. In particular, both the standard deviation of first-period consumption and the standard deviation of the fraction of domestic assets increase when one of the corner equilibria becomes unstable, remains relatively high during most of the motion to the opposite end of the parameter space and converges to zero when this new stationary solution is eventually approached. Note that this also implies that despite the near constancy of the mean value of $c_i(t)$, the dynamics is often characterized by a relatively wide range of individual choices. What happens is that after destabilization of an equilibrium, a broad range of choices of $c_i(t)$ and $f_i(t)$ gives higher utility (as can be inferred from the indifference curves in Fig. 5). Hence, many different types of mutations will be allowed to enter the population. The distribution of the choice variables spreads out and in the following, the whole population moves like a swarm from the left-hand side of the space of choice variables to its upper right-hand end (cf. Fig. 6). When the portfolio choices converge to a homogeneous situation $f_i(t) = 1$ again, higher indifference curves can only be reached with consumption levels close to the steady state level c^* . This leads to a decrease of the bandwidth of first-period consumption levels in the population. Eventually, the variability with respect to both choice variables shrinks to zero. Once all individuals are close to the utility-maximizing steady state levels ($f_i(t) = 1$, $c_i(t) = c^*$), a small

destabilizing mutation will lead to a sudden spread of the distribution of strategies and will set into motion a left-ward dynamics of the whole swarm of individuals.

Figs. 5 and 6 about here

7. Conclusions

Elaborating on the GA version of the Kareken-Wallace model introduced by Arifovic (1996) and Arifovic/Gencay (2000), we have analyzed both the potential and the limitations for this type of artificial open economy to generate realistic time series properties. As it turns out, the model can generate time series which very closely mimic the statistical characteristics of empirical data. The mechanism responsible for the emergence of these interesting dynamics seems to be similar to the one analyzed within a different context by Lux and Marchesi (1999): the model has a continuum of equilibria with an indeterminate distribution of strategies among agents (as has been argued above, any distribution of the f would be admissible in equilibrium). With the stochasticity of the genetic process, there will always be distortions preventing the system from settling at any particular equilibrium. Because of the evolutionary instability of any distribution of strategies these random distortions will evoke self-amplifying tendencies which produce large price changes (fat tails) and volatility clustering. However, we also find that a small probability of mutation and a small number of agents are needed to get this realistic output for the exchange rate. With a large population, the destabilizing tendencies are so strong that the crucial choice variable, f , bounces back and forth between the corners of the admissible parameter space. This applies to both binary and real coded GAs. While the requirement of small mutation rates might be considered to be plausible and not too restrictive, having to restrict the population size to numbers below, say, $N = 1000$ is much more cumbersome. Real markets (in particular, the world-wide market for foreign exchange), surely have more participants so that $N < 1000$ seems an unrealistic requirement. However, this disappointing finding is shared by other multi-agent models (cf. Egenter *et al.*, 1999, Yeh, 2001, Challet and Marsili, 2002). Essentially, with high N , a law of large numbers becomes effective even in models with a large number of available strategies and the randomness from the interaction between the microscopic choice of strategies vanishes. While in certain models, prices converge to fundamental values in the large economy limit (Egenter *et al.*, 1999), the absence of fundamentals in the Kareken-Wallace model appears to be responsible for the oscillations between extreme choices.

How could one overcome these uncomfortable findings and save the ‘nice’ results obtained with smaller populations? One possibility would be to allow for more coherence among individuals via social sharing of information. Allowing for groups of agents to form, we would get a smaller *effective* number of agents. As an alternative, endogenous development of wealth could lead to some agents exerting more influence on the market outcome than others (of course, this feature would be

particularly difficult to incorporate into the present simple model). This would presumably also change the outcome in a way that differs from the atomistic case analyzed above. Exploring these avenues is left for future research.

References:

- Arifovic, J., 1996, The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies, *Journal of Political Economy* 104, pp. 510 – 541.
- Arifovic, J. and R. Gencay, 2000, Statistical Properties of Genetic Learning in a Model of Exchange Rate, *Journal of Economic Dynamics and Control* 24, pp. 981 – 1005.
- Bornholdt, S., 2001, Expectation Bubbles in a Spin Model of Markets: Intermittency from Frustration Across Scales, *International Journal of Modern Physics C* 12, pp. 667 – 674.
- Challet, D. and M. Marsili, 2002, *Criticality and Finite-Size Effects in a Simple Realistic Model of Stock Market*. Mimeo: Oxford University (<http://xxx.lanl.gov/archive/cond-mat/0210549>).
- Chen, S.-H., T. Lux and M. Marchesi, 2001, Testing for Nonlinear Structure in an ‘Artificial’ Financial Market, *Journal of Economic Behavior and Organization*, pp. 327 – 342.
- Chiarella, C. and T. He, 2001, Asset Price and Wealth Dynamics under Heterogeneous Expectations, *Quantitative Finance* 1, pp. 509 - 526
- Dacorogna, M., R. Gencay, U. Müller, R. Olsen and O. Pictet, 2001, *An Introduction to High-Frequency Finance*. New York.
- Dawid, H., 1999, *Adaptive Learning by Genetic Algorithms: Analytical Results and Applications to Economic Models*. 2nd ed., Berlin: Springer.
- Egenter, E., T. Lux and D. Stauffer, 1999, Finite-Size Effects in Monte Carlo Simulations of Two Stock Market Models, *Physica A*, pp. 250 – 256.
- Eshelman, L. and J. Schaffer, 1993, Real-Coded Genetic Algorithms and Interval Schemata, in: *Foundations of Genetic Algorithms 2*. San Mateo: Morgan Kaufman Publ.
- Gaunersdorfer, A. and C. Hommes, 2000, *A Nonlinear Structural Model for Volatility Clustering*. Mimeo: University of Vienna.
- Geweke, J. and S. Porter-Hudak, 1983, The Estimation and Application of Long-Memory Time Series Models, *Journal of Time Series Analysis* 4, pp. 221 – 238.
- Herrera, F., M. Lozano and J. Verdegay, 1998, Tackling Real-Coded Genetic Algorithms: Operators and Tools for Behavioural Analysis, *Artificial Intelligence Review* 12, pp. 265 – 319.
- Holland, J., 1975, *Adaptation in Natural and Artificial Systems*. Ann Arbor: University of Michigan Press.
- Iori, G., 2002, A Microsimulation of Traders’ Activity in the Stock Market: the Role of Heterogeneity, Agents’ Interactions and Trade Frictions, *Journal of Economic Behavior and Organization* 49, 269 - 285
- Kareken, J. and N. Wallace, 1982, On the Indeterminacy of Equilibrium Exchange Rates, *Quarterly Journal of Economics* 96, pp. 207 – 222.

Kirman, A. and G. Teysiere, 2001, *Financial Markets with Interacting Agents: the Microeconomic Origins of Long Memory*. Mimeo: University Aix-Marseille.

Lux, T. and M. Marchesi, 1999, Scaling and Criticality in a Stochastic Multi-Agent Model of a Financial Market, *Nature* 397, 498 - 500.

Lux, T. and M. Marchesi, 2000, Volatility Clustering in Financial Markets: A Micro-Simulation of Interacting Agents, *International Journal of Theoretical and Applied Finance* 3, 675 – 702.

Meese, R. and K. Rogoff, 1983, Empirical Exchange Rate Models of the Seventies; Do the Fit Out of Sample, *Journal of International Economics* 14, 3 – 24.

Prügel-Bennett, A. and J. Shapiro, 1994, Analysis of Genetic Algorithms Using Statistical Mechanics, *Physical Review Letters* 72, 1305 – 1309.

Sargent, T., 1993, *Bounded Rationality in Macroeconomics*. Oxford: University Press.

Srinivas, M. and L. Patnaik, 1996, Genetic Search: Analysis Using Fitness Moments, *IEEE Transactions on Knowledge and Data Engineering* 8, 120 – 133.

Vries, C. de, 1994, Stylized Facts of Nominal Exchange Rates, in, F. van der Ploeg, ed., *The Handbook of International Macroeconomics*. Oxford: Blackwell Publ.

Yeh, C.-H., 2001, *The Influence of Market Size in an Artificial Stock Market: The Approach Based on Genetic Programming*. Mimeo: I-Shou University.

Table 1: Variation of Tail Index Estimate from Binary Coded GAs

Varying mutation probability, constant population size N= 60

P _{mut}	tail size 2.5 %			tail size 5%			tail size 10 %		
	min	median	max	min	median	max	min	median	max
0.005	1.37	2.82	5.07	1.20	2.16	4.81	1.15	1.92	5.37
0.010	1.93	3.36	5.06	1.35	2.81	3.81	1.34	2.21	3.06
0.015	2.16	3.90	5.87	1.76	3.25	4.38	1.60	2.55	3.36
0.020	1.88	4.14	6.30	2.15	3.51	4.80	2.13	2.83	3.75
0.025	2.60	4.45	6.87	2.60	3.81	5.47	2.31	3.04	3.72
0.030	3.25	4.75	7.00	3.24	4.11	5.27	2.79	3.34	3.90
0.035	2.30	5.15	7.64	2.52	4.25	5.34	2.80	3.46	4.01
0.040	3.50	5.13	8.38	3.23	4.42	6.64	2.85	3.61	4.40
0.045	3.50	5.09	7.37	3.48	4.46	6.36	3.04	3.71	4.52
0.050	3.88	5.20	7.62	3.56	4.53	5.70	3.10	3.76	5.01

Constant mutation probability p_{mut} = 0.01, varying population size

N	tail size 2.5 %			tail size 5%			tail size 10 %		
	min	median	max	min	median	max	min	median	max
20	1.73	3.29	5.43	1.38	2.64	4.67	1.30	2.26	9.17
100	2.76	3.77	4.56	2.12	3.17	3.85	2.02	2.49	2.79
200	2.29	4.19	5.67	2.26	3.54	4.64	2.16	2.68	3.16
1000	1.87	3.22	5.54	1.86	3.15	4.58	1.85	2.80	3.81
2000	1.46	2.64	4.70	1.60	2.74	4.17	1.78	2.67	3.74
4000	1.19	1.94	3.78	1.52	1.92	3.53	1.56	2.20	3.21
10000	1.44	1.92	3.27	1.28	1.82	3.22	1.35	1.93	2.84

Note: Table 1 shows estimates of the parameter α from the asymptotic power-law behavior of large returns:

$\Pr \text{ob.}(|r_t| > x) \sim x^{-\alpha}$. We follow the literature in applying a conditional maximum likelihood estimator with a prespecified size of the tail region. To explore the sensitivity of the tail index estimates with respect to the choice of the cut-off, we tried tail regions of 2.5%, 5% and 10%. Empirical estimates usually show a certain tendency of increasing tail indices when the tail size is reduced. For variation of the mutation probability, the minimum, median and maximum over 100 samples with 2,000 data points each are shown. For variation of the number of agents, only 25 samples were used due to the increase in computation time with increasing number of GA chromosomes.

Table 2: Variation of Index of Fractional Differentiation from Binary Coded GAsVarying mutation probability, constant population size $N=60$

p_{mut}	raw			squared			absolute returns		
	min	median	max	min	median	max	min	median	max
0.005	-1.05	-0.03	0.21	-0.23	0.25	0.76	-0.09	0.36	0.81
0.010	-0.95	-0.05	0.24	-0.04	0.31	0.77	0.15	0.42	0.78
0.015	-0.61	-0.12	0.29	0.02	0.34	0.81	0.11	0.43	0.88
0.020	-0.46	-0.17	0.07	-0.11	0.30	0.72	0.04	0.38	0.76
0.025	-0.57	-0.25	0.03	-0.08	0.25	0.82	-0.09	0.35	0.73
0.030	-0.58	-0.30	0.07	-0.11	0.23	0.51	-0.07	0.29	0.60
0.035	-0.66	-0.37	-0.07	-0.10	0.22	0.50	-0.09	0.28	0.47
0.040	-0.73	-0.43	-0.15	-0.06	0.19	0.44	-0.13	0.22	0.48
0.045	-0.76	-0.46	-0.14	-0.12	0.19	0.49	-0.07	0.22	0.50
0.050	-0.84	-0.51	-0.18	-0.13	0.15	0.47	-0.10	0.18	0.52

Constant mutation probability $p_{mut} = 0.01$, varying population size

N	raw			squared			absolute returns		
	min	median	max	min	median	max	min	median	max
20	-0.95	-0.02	0.32	-0.01	0.21	0.80	0.04	0.33	0.90
100	-0.31	-0.06	0.27	-0.03	0.30	0.79	0.10	0.40	0.73
200	-0.51	-0.18	0.04	0.09	0.30	0.79	0.20	0.35	0.66
1000	-0.76	-0.56	-0.32	-0.12	0.13	0.34	0.01	0.16	0.34
2000	-0.92	-0.61	-0.11	-0.15	0.05	0.27	-0.06	0.07	0.27
4000	-0.98	-0.48	-0.20	-0.24	0.04	0.33	-0.20	0.07	0.36
10000	-0.85	-0.40	-0.11	-0.23	0.01	0.13	-0.12	0.04	0.22

Note: Table 2 shows estimates of the parameter d from the hyperbolic decay of auto-covariances for variables with long-term dependence: $E[x_t x_{t-\Delta t}] \sim \Delta t^{2d-1}$. We estimate d via the log periodogram regression technique proposed by Geweke and Porter-Hudak (1983). The underlying data are the same as in Table 1. For variation of the mutation probability, the minimum, median and maximum over 100 samples with 2,000 data points each are shown. For variation of the number of agents, only 25 samples were used due to the increase in computation time with increasing number of GA chromosomes.

Table 3: Results of Unit-Root Tests for Binary Coded GAs

Varying mutation probability, constant population size N= 60, 100 runs								
P _{mut}	DF test			Rejections	ADF test			Rejections
	min	median	max		min	median	max	
0.005	0.15	0.97	1.00	89 (85)	0.30	0.99	1.00	54 (47)
0.010	0.47	0.97	1.00	97 (95)	0.51	0.98	1.03	64 (59)
0.015	0.23	0.96	0.98	100 (100)	0.55	0.97	0.99	95 (91)
0.020	0.11	0.95	0.98	100 (100)	0.42	0.96	0.99	99 (99)
0.025	0.61	0.93	0.97	100 (100)	0.79	0.95	0.98	100 (100)
0.030	0.68	0.92	0.96	100 (100)	0.70	0.94	0.98	100 (100)
0.035	0.20	0.91	0.96	100 (100)	0.44	0.93	0.96	100 (100)
0.040	0.49	0.91	0.95	100 (100)	0.64	0.93	0.96	100 (100)
0.045	0.27	0.90	0.95	100 (100)	0.37	0.92	0.95	100 (100)
0.050	0.15	0.90	0.97	100 (100)	0.33	0.91	1.01	99 (100)

Constant mutation probability, p _{mut} = 0.01, varying population size, 25 runs								
P _{mut}	DF test			Rej.	ADF test			Rej.
	min	median	max		min	median	max	
20	0.14	0.95	0.99	25 (25)	0.21	0.98	1.00	14 (12)
100	0.91	0.98	0.99	21 (21)	0.94	0.99	1.00	18 (18)
200	0.86	0.98	0.99	25 (24)	0.88	0.98	0.99	24 (23)
1000	0.32	0.92	0.98	25 (25)	0.41	0.91	0.97	25 (25)
2000	0.16	0.78	0.98	25 (25)	0.17	0.86	0.97	25 (25)
4000	0.13	0.77	0.96	25 (25)	0.15	0.81	0.97	25 (25)
10000	0.21	0.63	0.91	25 (25)	0.26	0.56	0.93	25 (25)

Note: Table 3 shows estimates of the parameter \tilde{n} from a regression of the log exchange rate on its lagged value. The columns labeled 'rejection' give the number of cases in which we can reject the unit root null hypothesis $\tilde{n} = 1$ from the one-sided (two-sided) DF and ADF tests. Again, the underlying data are the ones already analysed in Tables 1 and 2. Hence, for variation of the mutation probability, the minimum, median and maximum over 100 samples with 2,000 data points each are shown. For variation of the number of agents, only 25 samples were used due to the increase in computation time with increasing number of GA chromosomes.

Table 4: Variation of Tail Index Estimate from Real Code d GAsVarying mutation probability p_{mut} , constant mutation variance $\hat{\sigma}_{mut} = 0.1$ and population size $N = 100$

P_{mut}	Tail size 2.5 %			Tail size 5%			Tail size 10 %		
	min	median	max	min	median	max	min	median	max
0.01	1.33	2.53	4.75	1.46	2.51	3.85	1.49	2.26	2.94
0.02	1.75	2.57	3.93	1.62	2.33	3.47	1.54	2.08	3.18
0.03	0.81	3.01	6.89	0.85	2.57	3.33	0.67	2.24	2.70
0.04	0.70	2.94	4.70	0.74	2.45	3.62	0.89	2.08	3.07
0.05	0.96	3.23	5.02	0.73	2.68	3.62	0.85	2.33	3.09
0.06	0.85	2.33	3.09	0.80	2.67	4.24	1.00	2.22	3.44
0.07	0.91	2.98	5.98	1.11	2.60	4.60	0.71	2.24	3.29
0.08	1.03	2.93	6.18	1.04	2.62	4.65	0.86	2.14	3.51
0.09	0.99	2.72	6.44	1.03	2.35	4.63	0.93	2.02	3.67
0.10	1.07	2.47	6.08	0.95	2.13	3.99	0.97	1.84	3.60

Constant mutation probability $p_{mut} = 0.05$, constant population size $N = 100$, varying mutation variance

$\hat{\sigma}_{mut}$	Tail size 2.5 %			Tail size 5%			Tail size 10 %		
	min	median	max	min	median	max	min	median	max
0.025	1.40	3.61	5.52	1.04	3.12	4.07	1.04	2.60	3.10
0.050	1.37	3.46	5.56	0.92	3.00	4.32	0.77	2.54	3.38
0.075	0.67	3.37	6.34	0.65	2.95	3.97	0.59	2.47	3.32
0.100	0.78	3.28	5.03	0.79	2.77	4.08	1.00	2.27	3.07
0.125	0.67	3.05	4.67	0.74	2.61	3.81	0.75	2.25	3.08
0.150	0.70	2.91	5.15	0.89	2.57	4.41	0.71	2.14	2.99
0.175	0.90	2.86	5.27	0.95	2.53	3.96	0.95	2.12	3.08
0.200	0.95	2.85	4.78	0.95	2.40	3.88	0.97	1.98	3.04

Constant mutation probability $p_{mut} = 0.05$ and variance $\hat{\sigma}_{mut} = 0.025$, varying population size

N	Hill 2.5 %			Hill 5%			Hill 10 %		
	min	median	max	min	median	max	min	median	max
20	1.32	3.29	7.57	0.61	2.75	5.22	0.51	2.30	3.49
100	1.46	3.85	6.08	1.49	3.16	4.22	1.41	2.68	3.23
200	1.35	4.23	5.73	1.06	3.50	4.43	0.93	2.85	3.85
1000	0.77	4.22	8.13	0.71	3.40	6.11	0.74	3.00	4.43
2000	0.85	2.65	8.16	0.68	2.39	5.97	0.62	2.08	4.52
4000	0.70	1.56	10.04	0.66	1.37	7.20	0.59	1.18	5.06
10000	0.62	1.44	8.82	0.55	1.13	6.34	0.61	0.96	5.09
20000	0.52	1.49	7.99	0.62	1.14	7.24	0.58	0.86	5.04

Note: Table 4 shows estimates of the parameter $\hat{\alpha}$, now for real-coded GAs. Unlike in Figs. 1 to 3, we always give the minimum, median and maximum over 100 replications with 2,000 data points each even in the case of varying number of agents (real GAs are less demanding in terms of computation time than binary ones).

Table 5: Variation of Index of Fractional Differentiation from Real Coded GAsVarying mutation probability p_{mut} , constant mutation variance $\sigma_{mut} = 0.1$ and population size $N = 100$

P_{mut}	raw			squared			absolute returns		
	min	median	max	min	median	max	min	median	max
0.01	-1.06	-0.39	0.24	-0.28	0.09	0.41	-0.22	0.17	0.52
0.02	-0.79	-0.05	0.25	-0.16	0.14	0.43	-0.22	0.24	0.55
0.03	-0.44	-0.04	0.32	-0.14	0.13	0.90	-0.08	0.22	1.11
0.04	-0.52	-0.06	0.22	-0.08	0.28	1.04	-0.10	0.40	1.07
0.05	-0.44	-0.03	0.25	-0.11	0.25	1.17	-0.01	0.35	1.16
0.06	-0.48	-0.08	0.29	-0.02	0.28	1.09	-0.03	0.40	1.23
0.07	-0.52	-0.08	0.21	-0.11	0.30	0.87	-0.06	0.47	1.03
0.08	-0.43	-0.07	0.22	-0.03	0.34	0.93	-0.04	0.48	0.99
0.09	-0.52	-0.11	0.26	-0.01	0.34	0.89	-0.05	0.50	1.32
0.10	-0.48	-0.13	0.17	0.05	0.35	1.00	0.06	0.53	0.92

Constant mutation probability $p_{mut} = 0.05$, constant population size $N = 100$, varying mutation variance

σ_{mut}	raw			squared			absolute returns		
	min	median	max	min	median	max	min	median	max
0.025	-0.30	-0.01	0.34	-0.22	0.16	0.98	-0.10	0.24	0.98
0.050	-0.48	-0.07	0.21	-0.40	0.19	1.03	-0.06	0.26	1.01
0.075	-0.52	-0.04	0.30	-0.03	0.22	0.90	-0.04	0.29	1.14
0.100	-0.39	-0.05	0.27	-0.09	0.30	1.14	-0.07	0.45	1.08
0.125	-0.56	-0.06	0.31	-0.08	0.31	0.94	-0.18	0.46	1.03
0.150	-0.42	-0.06	0.17	-0.10	0.31	1.05	0.01	0.43	1.10
0.175	-0.50	-0.07	0.25	0.02	0.34	1.13	-0.06	0.49	1.07
0.200	-0.46	-0.09	0.19	-0.21	0.35	1.00	-0.02	0.50	1.00

Constant mutation probability $p_{mut} = 0.05$ and variance $\sigma_{mut} = 0.025$, varying population size

N	raw			squared			absolute returns		
	min	median	max	min	median	max	min	median	max
20	-0.48	-0.02	0.46	-0.10	0.21	0.74	-0.07	0.29	0.89
100	-0.37	0.00	0.21	-0.16	0.13	0.65	-0.16	0.20	0.71
200	-0.60	0.01	0.30	-0.16	0.12	0.73	-0.18	0.15	0.81
1000	-0.76	-0.02	0.40	-0.27	0.13	0.71	-0.26	0.22	1.05
2000	-0.66	0.03	0.39	-0.31	0.11	0.85	-0.33	0.28	1.02
4000	-0.79	-0.15	0.39	-0.31	0.10	0.39	-0.34	0.29	0.98
10000	-0.72	-0.28	0.45	-0.30	0.13	0.86	-0.28	0.37	0.96
20000	-0.75	-0.36	0.47	-0.13	0.10	0.86	-0.03	0.34	1.19

Note: Table 5 shows estimates of the parameter d , now for real-coded GAs. The data are the same as in Table 4.

Table 6: Results of Unit-Root Tests for Real Coded GAs

Varying mutation probability p_{mut} , constant mutation variance $\sigma_{mut} = 0.1$ and population size $N = 100$, 100 runs

p_{mut}	DF test			Rej.	ADF test			Rej.
	min	median	max		min	median	max	
0.01	0.29	0.90	1.00	74 (70)	0.42	0.94	1.00	59 (61)
0.02	0.83	0.99	1.00	54 (49)	0.87	0.99	1.00	54 (49)
0.03	-0.00	0.99	1.00	50 (47)	0.01	0.99	1.00	35 (34)
0.04	-0.00	0.99	1.00	49 (42)	0.06	0.99	1.00	34 (30)
0.05	0.01	0.99	1.01	46 (45)	0.26	0.99	1.00	39 (37)
0.06	-0.00	0.99	1.00	55 (45)	-0.00	0.99	1.00	40 (34)
0.07	0.00	0.99	1.00	56 (47)	0.00	0.99	1.01	40 (36)
0.08	0.00	0.99	1.00	58 (51)	0.01	0.99	1.00	46 (40)
0.09	-0.00	0.99	1.00	53 (48)	-0.00	0.99	1.00	48 (41)
0.10	-0.00	0.98	1.01	68 (66)	-0.00	0.98	1.01	68 (63)

Constant mutation probability $p_{mut} = 0.05$, constant population size $N = 100$, varying mutation variance, 100 runs

σ_{mut}	DF test			Rej.	ADF test			Rej.
	min	median	max		min	median	max	
0.025	-0.00	1.00	1.00	36 (30)	-0.00	1.00	1.00	25 (24)
0.050	0.03	0.99	1.00	43 (32)	0.19	1.00	1.00	26 (21)
0.075	0.01	0.99	1.00	51 (40)	0.01	0.99	1.00	34 (33)
0.100	-0.00	0.99	1.00	69 (60)	0.31	0.99	1.00	54 (48)
0.125	-0.00	0.99	1.00	68 (65)	-0.00	0.99	1.03	55 (53)
0.150	-0.00	0.99	1.00	64 (58)	-0.00	0.99	1.00	46 (39)
0.175	-0.00	0.99	1.00	76 (66)	0.12	0.99	1.10	50 (48)
0.200	-0.00	0.98	1.00	72 (69)	0.02	0.99	7.71	60 (57)

Constant mutation probability $p_{mut} = 0.05$ and variance $\sigma_{mut} = 0.025$, varying population size, 100 runs

N	DF test			Rej.	ADF test			Rej.
	min	median	max		min	median	max	
20	0.77	0.98	1.00	85 (79)	0.88	0.99	1.00	21 (17)
100	-0.00	1.00	1.00	29 (24)	0.00	1.00	1.00	19 (15)
200	0.99	1.00	1.00	2 (4)	0.99	1.00	1.00	4 (4)
1000	-0.00	1.00	1.76	30 (38)	-0.00	1.00	24.74	31 (33)
2000	-0.00	1.00	1.66	24 (33)	-0.00	1.00	6.11	27 (29)
4000	-0.00	1.00	1.06	37 (50)	-0.00	1.00	1.02	35 (43)
10000	0.01	1.00	1.52	45 (63)	0.01	1.00	12.52	44 (52)
20000	0.12	1.00	1.03	43 (54)	0.12	1.00	1.01	42 (49)

Note: Table 6 shows results from unit-root tests, now for real-coded GAs. The underlying data are the same as in Tables 4 and 5.




Fig. 1.: A typical ‘realistic’ series of returns from a simulated economy with a binary-coded GA population of 100 agents. For economic parameters, see main text. GA parameters are: $p_{\text{mut}} = 0.01$ and $p_{\text{cross}} = 0.6$.




Fig. 2.: Log exchange rates from simulated economies with $p_{\text{mut}} = 0.01$ (upper panel) and $p_{\text{mut}} = 0.05$ (lower panel). The population consists of 100 binary-coded GAs in both cases. Although not fully realistic, for parts of the upper time series the DF and ADF tests are unable to reject the unit root null. The simulation in the lower panel has more easily recognizable mean-reverting features.




Fig. 3.: Log exchange rate (top), average portfolio fraction of home currency (middle), and average first-period consumption (bottom) for a real-coded GA population of 200 agents. For economic parameters, see main text. GA parameters are: $p_{mut} = 0.05$, $\sigma_{mut} = 0.025$ and $\gamma = 0.2$. The dynamics seems to be characterized by unsystematic changes of the portfolio composition which lead to exchange rate fluctuations, but leave average consumption choices almost unaffected.




Fig. 4.: Log exchange rate (top), average portfolio fraction of home currency (middle), and average first-period consumption (bottom) for a real-coded GA population of 20,000 agents. For economic parameters, see main text. GA parameters are: $p_{\text{mut}} = 0.3$, $\sigma_{\text{mut}} = 0.025$ and $\gamma = 0.2$. The middle and bottom panels show both the mean (solid lines) and standard deviations (broken lines) of the distribution of the choice variables within the population. In order to lodge the mean and standard deviation in the bottom panel, we have subtracted the steady state value $c^* = 7$, reduced the standard deviation of $c_i(t)$ by one-half and magnified the standard deviation of $f_i(t)$ by a factor 10. The higher p_{mut} compared to Fig. 3 mainly serves to decrease the amplitude of the cycle and enhance convergence to regular cycles,




Fig. 5.: Indifference curves. The underlying utility function is $U = c(t) c(t+1)$, endowments are $w_1 = 10$, $w_2 = 4$. For this illustration it has been assumed that all agents have chosen their first-period consumption level equal to its steady state value, $c^* = 7$ and that the fraction of domestic money in the portfolios of generations t and $t+1$ has mean values $\overline{f(t+1)} = 0.55$ and $\overline{f(t)} = 0.5$.




Fig. 6.: A snapshot of the evolution of the population corresponding to one upward half-cycle in Fig. 4. The graph shows on its left-most part the distribution of choice parameters within a generation shortly after the lower turning point (triangles). The pluses and diamonds show the distribution of choice parameters within the same dynasty after 40 and 80 periods, respectively.