Computational Economics **21:** 209–229, 2003.

Modeling Exchange Rate Behavior with a Genetic Algorithm ⋆

C. LAWRENZ[1] and F. WESTERHOFF[2]

[1]*WestLB, Germany;* [2]*University of Osnabrueck, Department of Economics, Rolandstrasse 8, D-49069 Osnabrueck, Germany; E-mail: fwesterho@oec.uni-osnabrueck.de*

Abstract. Motivated by empirical evidence, we construct a model where heterogeneous, boundedly-rational market participants rely on a mix of technical and fundamental trading rules. The rules are applied according to a weighting scheme. Traders evaluate and update their mix of rules by genetic algorithm learning. Even for fundamental shocks with a low probability, the interaction between the traders produces a complex behavior of exchange rates. Our model simultaneously produces several stylized facts like high volatility, unit roots in the exchange rates, a fuzzy relationship between news and exchange-rate movements, cointegration between the exchange rate and its fundamental value, fat tails for returns, a declining kurtosis under time aggregation, weak evidence of mean reversion, and strong evidence of clustering in both volatility and trading volume.

Key words: exchange rate theory, technical and fundamental trading rules, genetic algorithm

1. Introduction

Since the development of real-time information systems and the decline of transaction costs following the liberalization of capital markets in the mid 80s, both daily foreign- exchange turn-over and the volatility in exchange rates have increased sharply. More and more the trading volume reflects very short-term transactions, indicating a highly speculative component (BIS, 1999). Surprisingly, when the market participants determine their speculative investment positions, rather simple technical and fundamental trading strategies are applied (Taylor and Allen, 1992).

The chartists-fundamentalists approach is a research direction that focuses on explaining such speculative transactions (Frankel and Froot, 1986; Kirman, 1993; Brock and Hommes, 1997). Of crucial importance in this class of models is the behavior of the so-called chartists and fundamentalists, because the interaction between these two groups has the potential for generating interesting non-linear dynamics. Recently, some multi-agent models in the spirit of the

⋆ Presented at the 6th International Conference of the Society for Computational Economics on Computing in Economics and Finance, July 6–8, 2000, Universitat Pompeu Fabra, Barcelona, Spain. We thank the participants and especially David Belsley and an anonymous referee for helpful comments.

chartists-fundamentalists models have emerged (LeBaron, 2000). Since these models allow for many interacting heterogeneous agents, the exchange-rate behavior becomes even more elaborate. The significance of these models is based upon their ability to match basic stylized facts of the empirical data.

LeBaron (2000) argues that multi-agent models have advantages and disadvantages. On the one hand, these artificial markets allow agents to explore a fairly wide range of possible forecasting rules. On the other hand, due to the inherent complexity it is often very difficult to pin down the causalities acting inside the market. Our paper tries to alleviate this criticism by following a suggestion of Mandelbrot (1997). According to Mandelbrot, a good model of price variation is one that mimics a great number of empirical facts within a simple framework.

Hence, this paper aims to develop a simple exchange-rate model in the vein of multi-agent systems to get a deeper understanding of the forces that drive foreign-exchange markets. Rather than deriving results from a well-defined utility maximization problem, details are used from the market microstructure and psychological evidence. We construct a model with heterogeneous, boundedly rational market participants relying on a mix of technical and fundamental trading rules. The rules are applied according to a weighting scheme. Traders evaluate and update their mix of rules from time to time. The selection process is modeled by a genetic algorithm, which has proven to be a useful tool for describing learning behavior (Dawid, 1999).

Our main result is that the interaction of the trading rules generates a complex behavior for exchange rates. For instance, the simulated time series resemble a stochastic trend (unit roots) in the first moments. Further, simulations of the model produce high volatility. Even though the relationship between news and exchange-rate movements is fuzzy in the short run, the exchange rate and its fundamental are cointegrated over longer time periods. The returns of the generated exchange rates show a high kurtosis that declines under time aggregation. Fat tails are also identified by the scaling behavior of the returns, which roughly follow a power law. In addition, weak evidence of short-run mean reversion and strong evidence of clustering in both volatility and trading volume are found. All these features are typically observed in the foreign exchange market.

The paper is organized as follows: Section 2 presents a genetic-algorithm exchange-rate model. Section 3 discusses some simulation results. The last section offers conclusions and possible extensions.

2. A Genetic Algorithm Exchange Rate Model

2.1. DESCRIPTION OF THE MARKET

The starting point for our approach is the observation that most of the trading volume in the foreign-exchange market is due to short-term speculative trading. To model the foreign-exchange market in a realistic perspective, we use empirical evidence to describe the speculative behavior of the traders.

The story of our model is as follows: The foreign-exchange market is ruled by a limited number of large traders (BIS, 1999). Each trader has a collection of strategies to determine his speculative investment positions. The set of applied rules is limited and common knowledge. The rules can be subsumed under technical and fundamental analysis (Taylor and Allen, 1992). Typically, they persist for a long period of time, just like the famous Dow-theory. However, the agents may alter the importance of each of the rules they follow. The decision about the mix of rules is made at the beginning of each trading period.

The evolution of the agents' mix of rules depends on a feedback process. We use a genetic algorithm to describe the social learning processes for the trader population. Learning takes place according to three different learning schemes. First, agents whose strategies lead to relatively poor performance (profit) give up their trading strategies and copy the strategy of a more successful market participant. This is called learning by imitation. Second, agents meet, talk to each other about their trading strategies and thus possibly exchange parts of each other's behavior. This can be interpreted as a proxy for learning by communication. Third, agents may learn by experimentation. That is, the agents change their own trading strategies slightly. Through this, new mixes of trading strategies emerge. Since the agents are boundedly rational, mistakes in the learning process might also occur.

Such agents' behavior is not irrational: Heiner (1983) argues that using simple behavioral rules can result from the uncertainty of distinguishing between preferred and less-preferred options. For example, the complexity inherent in the foreign-exchange market can cause each agent to have a gap between his competence in making optimal decisions and the actual difficulty of making the decisions. The wider the gap, the more likely the agents follow a rule-governed strategy. Thus, agents could not do much better than by following some adaptive scheme of behavior.

Compared with the amount of speculative trading, firms' transactions relating to international trade and risk management play a less important role (BIS, 1999). Since our focus is on modeling the interaction among the speculators, such activity is assumed to be randomly distributed. Every period, the exchange rate is determined by the market-clearing condition.

The remainder of this section formalizes the technical and fundamental trading rules. For simplicity, only three technical and three fundamental trading rules are allowed. Descriptions of the process for the arrival of news and the perceptions of traders are given, and, after solving the model, the learning schemes are introduced.

2.2. TECHNICAL TRADING RULES

Technical analysis is a trading method that attempts to identify trends and their reversals by inferring future price movements from those of the recent past.[1] We shortly introduce three technical trading rules. One of the rules is believed to identify trends, another to spot trend reversals, and the third combines the two.

A broader discussion of these rules is found in Murphy (1999), which is an often-cited manual of technical analysis. Note that technical analysis is a very common method to determine investment positions. As reported by Taylor and Allen (1992), most foreign-exchange dealers place at least some weight on technical analysis.

Simple technical trading rules use only past movements in the exchange rate S as an indicator of market sentiment to identify trends. The most popular technical trading rule for trend analysis is the simple moving average. For example, demand in period t might be expressed as

$$d_t^{C,1} = \alpha^{C,1}[0.6(LogS_{t-1} - LogS_{t-2}) + 0.4(LogS_{t-2} - LogS_{t-3})]\,. \tag{1}$$

Equation (1) captures the typical behavior of the chartists. In general, chartists buy (sell) foreign currency if the exchange rate rises (declines). This rule is only applied if the actual exchange rate trend exceeds a certain threshold, say if $|(S_{t-1} - S_{t-2})/S_{t-2}| > 0.005$. Under this rule, if the trend is too weak (less than half a percent), the demand is zero. Since traders pay more attention to the most recent trend, a larger coefficient is selected for the first extrapolating than for the second term (0.6 versus 0.4). The coefficient $\alpha^{C,1}$ calibrates the relationship of the demand under this rule to that under the other rules.

To use the chartists' language, however, the simple moving-average rule is a follower, not a leader. It does not anticipate, but merely reacts, to the dynamics. The agents try to overcome this disadvantage by combining short-run and long-run moving averages. For example, the so called double-crossover method produces a buy signal if the shorter average crosses above the longer. The demand from this kind of rule could be formalized as

$$\begin{aligned} d_t^{C,2} \; = \; & \alpha^{C,2}[(LogS_{t-1} - LogS_{t-2}) - \\ & 0.5((LogS_{t-1} - LogS_{t-2}) + (LogS_{t-2} - LogS_{t-3}))]\,, \end{aligned} \tag{2}$$

where $\alpha^{C,2}$ is again a coefficient to adjust the demand. The first bracket comprises the fast-moving average and the second bracket the slowmoving average. With this rule, a trading signal is generated if the actual trend of the exchange rate falls behind the long-term trend.

When the market is not trending, because the price fluctuates in a horizontal band, chartists often rely on oscillator techniques to spot trend reversals. Through these rules, overbought and oversold conditions of a market are indicated. For instance, a market is said to be overbought when it is near an upper extreme. Thus, a warning signal is given that the price trend is overextended and vulnerable. Momentum rules are the most popular application of oscillator analyses. Trading signals are derived by comparing the velocity (momentum) of price changes. Demand from these rules might be written as

$$d_t^{C,3} = \alpha^{C,3}[(LogS_{t-1} - LogS_{t-4}) - (LogS_{t-2} - LogS_{t-5})]\,. \tag{3}$$

This demand function states that the chartists expect a future increase in the exchange rate when the observed change between period $t - 1$ and $t - 4$ relative to the change from period $t - 2$ and $t - 5$ starts increasing.

Note that by (1), (2) and (3), chartists place a market order today in response to past price changes, but price changes between period t and $t-1$ are disregarded. Such a lag structure is typical for technical trading rules, because only the past movements of the exchange rates are taken into account (Murphy, 1999).

2.3. FUNDAMENTAL TRADING RULES

Fundamental trading rules are based on the premise that the exchange rate converges towards its fundamental value. They deliver a buy (sell) signal, if the expected future exchange rate is above (below) the spot rate. However, the expected adjustment process might be formulated in different ways. We allow for three different kinds of expectation formation.[2]

The first specification of the expectation formation process of the fundamentalists is modeled in the classical regressive manner: that is, if an exchange rate deviates from its perceived equilibrium value S^{FP}, it is expected to return. Hence, we assume $E_t[S_{t+1}] = \beta S^{FP}_{t-1} + (1-\beta)S_{t-1}$, where β is the expected adjustment speed of the exchange rate towards its fundamental. Since the expectation formation for the trading period t has to be made in advance, the last available fundamental value is from $t-1$. The demand of fundamentalists might therefore be written as

$$\begin{aligned} d^{F,1} &= \alpha^{F,1}(E^{F,1}_t[S_{t+1}] - S_t)/S_t \\ &= \alpha^{F,1}(0.85 S^{FP}_{t-1} + 0.15 S_{t-1} - S_t)/S_t \,, \end{aligned} \tag{4}$$

where the demand depends on the relative distance between the expected and spot rates, and on the reaction coefficient $a^{F,1}$. We assume 0 to be 0.85, that is, the agents expect an adjustment of 85% of the spot rate towards its fundamental.

The second specification is a variation of the first. The expectation formation is not only regressive, but also incorporates an extrapolative component: $E_t[S_{t+1}] = \beta_1 S^{FP}_{t-1} + (1-\beta_1)(S_{t-1} + \beta_2(S_{t-1} - S_{t-2}))$, with $\beta_1 = \beta_2 = 0.5$. Thus, an agent using this rule expects an adjustment of the exchange rate towards its fundamental of 50% but corrects the speed of adjustment by the most recent exchange-rate movement. The agent takes into account that the market dynamics are influenced by technical trading rules. The demand that results is

$$\begin{aligned} d^{F,2}_t &= \alpha^{F,2}(E^{F,2}_t[S_{t+1}] - S_t)/S_t \\ &= \alpha^{F,2}[0.5 S^{FP}_{t-1} + 0.5(S_{t-1} + 0.5(S_{t-1} - S_{t-2})) - S_t]/S_t \,. \end{aligned} \tag{5}$$

In the case of larger exchange-rate movements, the extrapolating term may even outpace the expected adjustment.

When the agents are uncertain about the fundamental exchange rate, they allow themselves to be guided by past values when forming new expectations. These past values act as anchors for individual judgments about the future exchange rate. This phenomenon is called anchoring heuristics. In such periods, the formation of exchange-rate expectations is not only regressive but also anchored to the last few observations of the exchange rate. As in the case of (5), the importance of the fundamental exchange rate is lower than in the first expectation formation hypothesis. Assuming $E_t[S_{t+1}] = 0.15S^{FP}_{t-1} + 0.425(S_{t-1} + S_{t-2})$, demand becomes

$$\begin{aligned} d^{F,3}_t &= \alpha^{F,3}(E^{F,3}_t[S_{t+1}] - S_t)/S_t \\ &= \alpha^{F,3}[0.15S^{FP}_{t-1} + 0.425(S_{t-1} + S_{t-2}) - S_t]/S_t\,, \end{aligned} \tag{6}$$

where the fundamentalists now use the exchange rate in $t-1$ and $t-2$ as a basis for expectation formation.

2.4. NEWS ARRIVAL PROCESS

Implicitly, we assume that the traders form their expectations of the fundamental exchange rate on the basis of a structural model that is correct on average. The perception of the fundamental value is due to the news arrival process and behaves like a jump process. The logarithm of S^F is given by

$$LogS^F_t = LogS^F_{f-1} + p\varepsilon_t\,, \tag{7}$$

where the news ε_t is i.i.d. Normal with mean zero and (time invariant) variance σ^2. The random variable p is 1 with probability 0.2 and 0 otherwise. Thus, a shock hits the market on average every 5 periods. Although the agents follow the news arrival process very closely, they also commit temporary mistakes in news cognition. We assume

$$LogS^{FP}_t = LogS^F_t + \delta_t\,, \tag{8}$$

where the mistakes δ_t are i.i.d. Normally distributed with mean zero and constant variance.

2.5. ADDITIONAL TRANSACTIONS

The volume of transactions resulting from international trade and risk management is considerably smaller than that of speculative trading positions (BIS, 1999). Since we concentrate on speculative trading, we assume such demand to be randomly distributed. The demand by the firms is

$$d^T_t = \chi_t\,, \tag{9}$$

where χ_t is i.i.d. Normally distributed with mean zero and constant variance.

2.6. SOLUTION OF THE MODEL

The market-clearing condition is given as a sum over all trading positions of the traders and the excess demand of the firms; thus

$$\sum_{j=1}^{3}\sum_{i=1}^{N}\gamma^{C,i,j}d_t^{C,j}+\sum_{j=1}^{3}\sum_{i=1}^{N}\gamma^{F,i,j}d_t^{F,j}+d_j^T=0\,, \tag{10}$$

where N is the number of traders and γ is the weight of the technical or fundamental trading rule j of trader i. Using (4) to (6) and solving (9) for the exchange rate yields

$$S_t=\frac{\sum_{j=1}^{3}\sum_{i=1}^{N}\gamma^{F,i,j}\alpha^{F,j}E_t^{F,j}[S_{t+1}]}{\sum_{j=1}^{3}\sum_{i=1}^{N}\gamma^{F,i,j}\alpha^{F,j}-\sum_{j=1}^{3}\sum_{i=1}^{N}\gamma^{C,i,j}d_t^{C,j}-d_t^T}\,. \tag{11}$$

Since (11) precludes closed analysis, we simulate the dynamics to demonstrate that the underlying structure gives rise to complex exchangerate behavior, as is observed empirically.[3]

2.7. GENETIC ALGORITHM LEARNING

Finally, we must model the evolution of the weighting schemes of the traders. We assume, the market is established by $N = 30$ traders. Before the beginning of a trading period the trader fixes the technical and fundamental fraction of his demand by attaching a weight to every rule (Figure 1). Through this, the demand resulting from the technical rules and the total weight of the fundamentalists is determined. Using the market-clearing condition (10), the exchange rate in period t equalizes the fundamentalists demand on the one side and the chartists' and the firms' excess demand on the other. After the revelation of any news, the success of the mix of the rules is evaluated. Each trader i is totally defined by a string $p_t^{i,j}$ of real numbers of length $l = 8$. The first two elements encode the total weight of the fundamental and the technical trading rules, respectively. The elements $\rho_t^{i,3}$ to $\rho_t^{i,8}$ define the weight of the three fundamental and the three technical rules. The weight, for example, of the momentum rule C3 is specified by

$$\gamma_t^{C,i,3}=\rho_t^{i,2}\cdot\rho_t^{i,8}\,, \tag{12}$$

with $p_t^{i,1}+\rho_t^{i,2}=1$ and $\sum\rho_t^{i,j}=1$. Thus, the momentum rule contributes $\gamma_t^{C,i,3}\cdot d_t^{C,3}$ to the total demand expressed by trader i.

After the exchange rate and the fundamental value of the exchange rate have been revealed, the traders assess their success relative to the other traders (fitness).

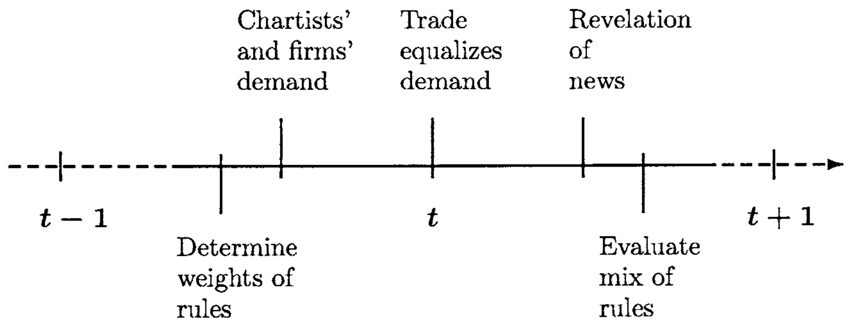

Figure 1. Time structure of the model.

The fitness is defined by the profit over the last five periods and the actual value of the stock:[4]

$$fitness(t, i) = \sum_{k=0}^{4}(-d^i_{t-k} \cdot S_{t-k}) + \sum_{k=0}^{4} d^i_{t-k} \cdot S_t \cdot r_t \,, \tag{13}$$

where

$$\begin{aligned} r_t \;=\; & 1 - (4{,}000(Log_{10}S_t - Log_{10}S^F_t)^3 + \\ & 0.5(Log_{10}S_t - Log_{10}S^F_t))/10 \,. \end{aligned} \tag{14}$$

The stock is valued with the actual exchange rate S_t weighted by an S-shaped risk function r_t that is restricted to the interval [0.5, 1.5]. This parameter represents the anticipation of the riskiness of the investment positions by the traders. If the exchange rate is above (below) its fundamental, r_t is smaller (larger) than 1. The fitness values are processed by the selection procedure. Through selection, the mixes with high fitness are joined to build up the material for the next generation. To prevent overselection and convergence in the early periods, we realize a tournament selection (Mitchell, 1996). Every trader chooses randomly a certain number of mixes from the current population. With probability $p_{sel} = 0.75$ he copies the one with the higher fitness, thus imitating successful behavior, and with a probability of 0.25, he sticks to his old rule mix.

Afterwards, a simple one-point crossover takes place. With probability $p_{cross} = 0.7$, a trader communicates; that is, he exchanges parts of his information with neighboring traders and incorporates rule mixes that have performed well in other contexts. In a wider sense, this may be seen as a proxy for learning by communication.

The traders evaluate the success of their rules every five periods and adjust them through imitation and communication. However, after every trading period, mutation takes place to produce a new generation of rule mixes. Due to coding with real numbers, instead of the standard genetic-mutation operator, mutation takes place with probability $p_{mut} = 0.04$, which results in a uniformly distributed change of [–0.5, 0.5]. A trader who is uninfluenced by other traders' changes his rule mixes

Table I. Parameter settings.

Description of parameter	Symbol	Value
Reaction coefficient for C1	$\alpha^{C,1}$	1.1
Reaction coefficient for C2	$\alpha^{C,2}$	2.75
Reaction coefficient for C3	$\alpha^{C,3}$	0.55
Reaction coefficient for F1	$\alpha^{F,1}$	0.3
Reaction coefficient for F2	$\alpha^{F,2}$	0.55
Reaction coefficient for F3	$\alpha^{F,3}$	0.2
Distribution of news	ε	N (0, 0.00005625)
Probability of news	$prob(p = 1)$	0.2
Misperception	δ	N (0, 0.00005625)
Number of traders	N	30
Probability of imitation	p_{sel}	0.75
Probability of crossover	p_{cross}	0.7
Probability of mutation	p_{mut}	0.04
Distribution of demand of firms	χ	N (0, 0.00025)

randomly either by mistake or because he tries to discover new, better strategies on his own. The update is completed by normalizing the weights, summing up to 1, and the mix is executed in the next trading period.

3. Simulation Results

3.1. PARAMETER SPECIFICATION

In this section, we discuss the simulation results. The main parameters are displayed in Table I. Unfortunately, most of the parameters cannot be observed directly. They are chosen to match some basic empirical facts. Thus, we set the relation between the sum of the α^F and the sum of the α^C so that the returns (log of price changes) remain principally in the region of 2%. Each of the α^F and α^C are selected in order to ensure that the demand from each of the trading rules contributes significantly to the transactions. In comparison, the demand by the firms is roughly one third of the total (BIS, 1999). The development of the fundamental exchange rate accounts for the fact that fundamental shocks do not occur every period. We assume that there is on average one fundamental news release per week. The total weight of the technical trading rules is limited for every agent to [0.25, 0.75]. As reported by Taylor and Allen (1992), only a minor fraction of professional traders consider technical and fundamental analysis to be competing

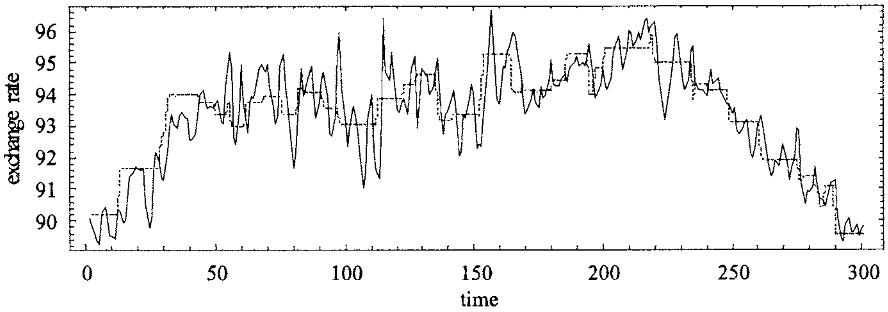

Figure 2. Simulated exchange-rate behavior for 300 periods. The black line is the exchange rate, the dashed line represents the fundamental exchange rate, $T = 300$ observations starting from $t = 4{,}650$.

to the point of being mutually exclusive. In general, both kinds of analysis are regarded as complementary. The parameters concerning the genetic algorithm are selected according to typical recommendations (Grefenstette, 1986).

3.2. EXCHANGE-RATE DYNAMICS

Figures 2 and 3 give a first impression of the resulting dynamics. Figure 2 contains the simulated dynamics for the exchange rate (solid line) and the stochastic development of its fundamental (dashed line) for 300 periods starting from period 4,650. Even a low probability of fundamental shocks has sufficed to generate complex exchange-rate movements, where the exchange rate fluctuates around the perceived fundamental exchange rate. For some time the exchange rate stays close to its fundamental, but this may change abruptly. Periods of lower volatility alternate with periods of higher volatility. Furthermore, the volatility of the exchange rate is far greater than the volatility of its fundamental.

Simplifying, the dynamics might be explained as follows. Technical trading rules always produce some kind of buy or sell signal. On the basis of a feedback process, a reinforcing run might emerge. But such a run cannot last because investment rules based on fundamentals work like a center of gravity. The more the exchange rate departs from the fundamental exchange rate, the higher the demand of the fundamentalists, until eventually their increasing net position triggers a mean reversion. But this indicates a new signal for the chartists and directly leads to the next phase. Extreme outliers can occur when the chartists have a clear trading signal and the influence of the fundamentalists is low. Most of the time, the dynamics are governed by the speculative behavior of the agents. Only in calm periods does the random demand of the firms have a more pronounced impact on exchange-rate fluctuations. Thus, the volatility of the foreign-exchange market need not be caused solely by exogenous shocks; it might be explained at least partially endogenously. The trading signals needed to keep the process going are generated by the agents themselves. This is exactly what Black (1986) calls noise trading. Black concludes

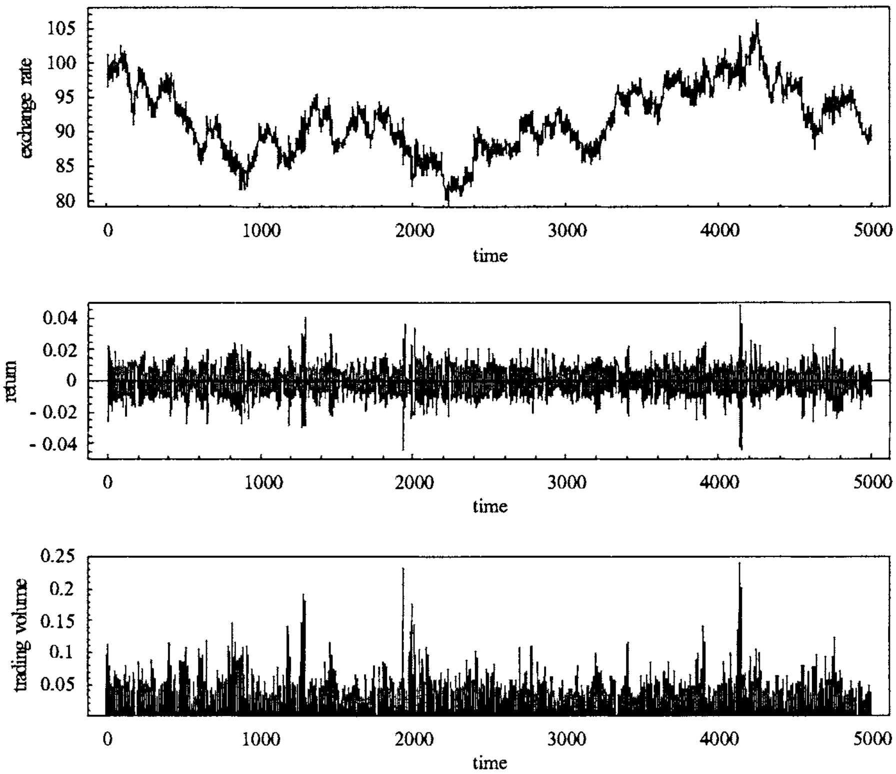

Figure 3. Exchange rates, returns, and trading volume for 5,000 periods. The top shows the exchange rate, the middle the returns (logarithm of price changes), and the bottom the trading volume, $T = 5{,}000$ observations.

that the noise of a large number of small events is essential to the existence of liquid markets. The argument is that a person who wants to trade needs another person with opposite beliefs. To explain the high trading volume in the foreign-exchange market, it is not reasonable to assume that differences in beliefs are merely the result of differences in information. In our model, noise (trading signals) is permanently produced by the agents.

Figure 3 contains the exchange-rate dynamics together with the corresponding development of the returns and the trading volume for 5,000 periods. The exchange rate resembles a stochastic trend. There are no clear patterns visible. The course of the returns demonstrates that deviations from the last rate occur mainly in a band of roughly 4%. Large returns coincide with large returns in the following periods, small returns with future small returns. The trading volume shows a similar clustering. These features are investigated in more detail later.

A well-known stylized fact is that exchange-rate time series have a unit root (Goodhart et al., 1993). The Augmented Dickey-Fuller (ADF) unit root testing procedure (Dickey and Fuller, 1981) tests the null hypothesis of difference stationary against the trend-stationary alternative. In particular, one estimates the regression

$$\Delta S_t = a_0 S_{t-1} + \sum_{j=1}^{l} a_j S_{t-j} , \tag{15}$$

where S denotes the exchange rate and Δ the first difference. The null hypothesis of a unit root is rejected if a_0 is negative and significantly different from zero. The lag length l can be chosen using data dependent methods. We find that for various lag settings and simulation runs the null hypothesis cannot be rejected. This result is often interpreted as evidence of a random walk behavior of the exchange rate (Brownian motion).

The top of Figure 4 displays the development of the strength of the three technical trading rules over 5,000 periods; the bottom shows the demand generated by the rules. The total weight of the chartists' positions varies between 25 and 55%. Most of the time they have less influence than that of the fundamental rules. Although the simple moving-average method is the dominant technical trading rule, for some periods the other two technical trading rules become leaders. Depending on the market circumstances, the profitability of the rules change. Because of the internal inertia of their adaptation process and the constrained optimization, the population sticks to the average weighting scheme to some degree. For instance, the periods 1,260 to 1,300 are characterized by a fast zigzag-course with frequent overshooting of the fundamental. In this environment the anchoring heuristic (F3) is the most appropriate strategy among the fundamental rules because it puts the least weight on the fundamental. But it takes up to period 1,290 before this rule replaces the double-crossover method as the most popular rule. Its average weight increases from 5% in period 1,250 to almost 70% in period 1,300. But even when the other fundamental rules become more profitable, it requires 40 periods before the importance of the anchoring heuristic falls again below the weight of the other fundamental rules. These waves of popularity can be observed throughout the entire course.

The bottom shows the demand generated by the technical trading rule for 300 periods. Each of the rules contributes significantly to the total transactions. The demand of the firms behaves randomly whereas the demand of the trading rules is correlated.

Figure 5 presents the corresponding runs for the fundamental trading rules. The influence of the fundamentalists is concentrated in the band from 45 to 75%. The strength of the fundamental trading rules is well mixed; each of the three fundamental trading rules has a period of dominance. Such a mix of trading strategies is close to what is reported in survey studies (Taylor and Allen, 1992).

The development of the power of the trading strategies in the time domain helps understanding the complicated exchange-rate dynamics. The exchange-rate fluctuations plotted in the Figures 2 and 3 can be classified into different regimes. Tranquil periods, where the spot rate is close to its fundamental, alternate with extremely volatile periods. A similar phenomenon is observed in (Lux and Marchesi,

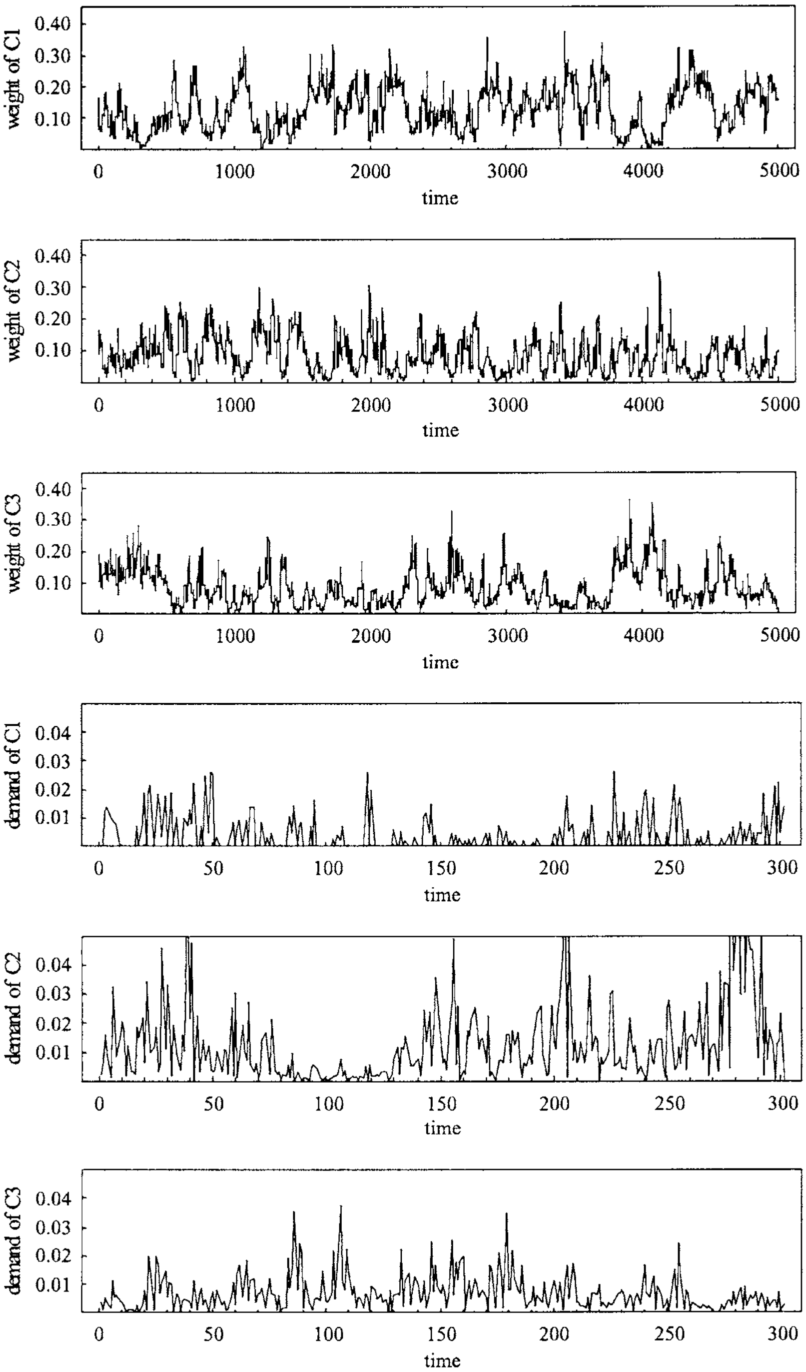

Figure 4. Weights and demand of the technical trading rules. The weights (demand) of the technical trading rules C1, C2, and C3 for $T = 5{,}000$ observations ($T = 300$ observations, starting from $t = 2{,}500$).

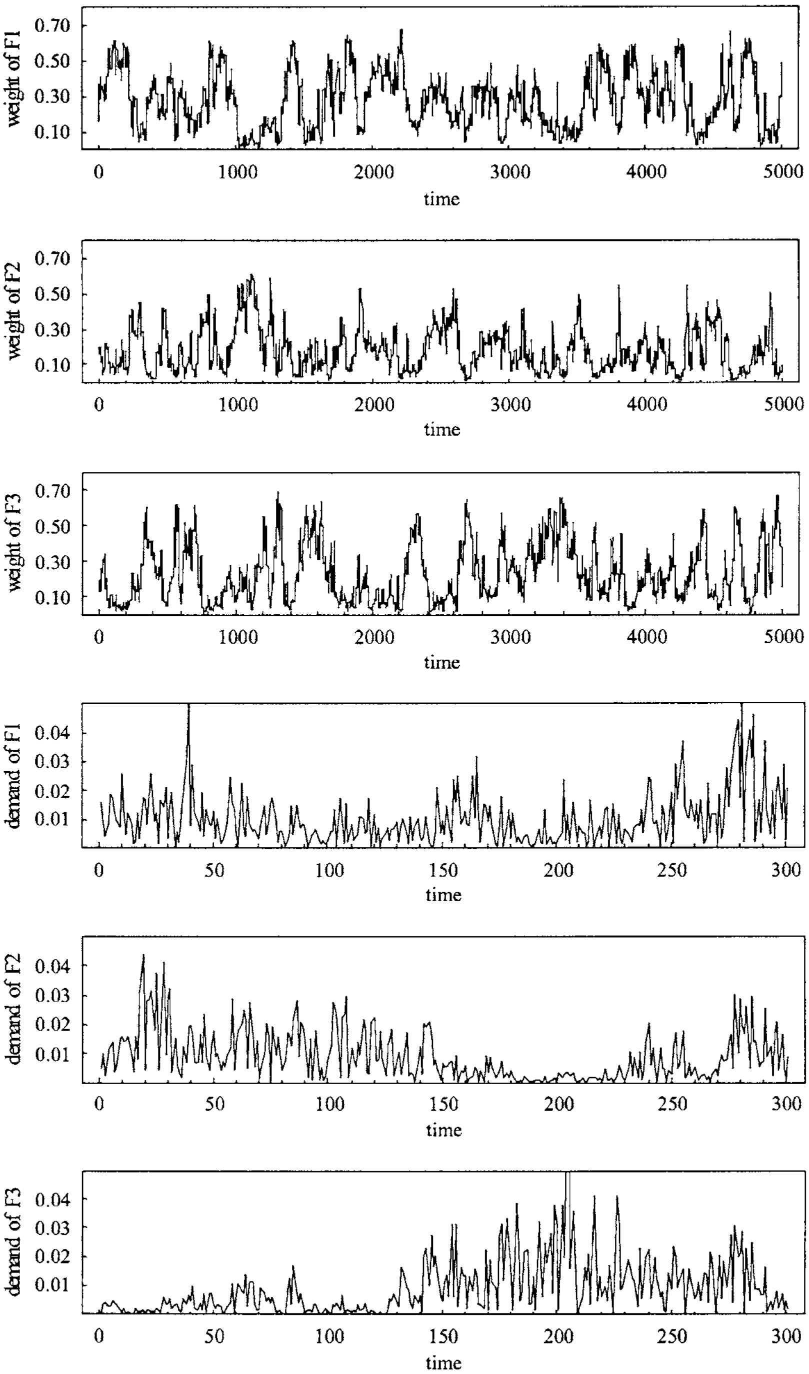

Figure 5. Weights and demand of the fundamental trading rules. The weights (demand) of the fundamental trading rules Fl, F2, and F3 for $T = 5{,}000$ observations ($T = 300$ observations, starting from $t = 2{,}500$).

2000). Also, longer swings in the exchange rate are followed by a faster zigzag-course. These different regimes result from the development of the mix of the trading rules.

Most importantly, the dynamics become less stable when the weight of chartists increases. For example, in period 4,130 to 4,150 a rise in the weight of the double-crossover method (C2) and the momentum rule (C3) triggers extreme movements of the exchange rate (Figures 2 and 3). They generate a self-fulfilling trend forecast with a highly volatile sequence and a large increase in trading volume – although almost no change in the fundamental value takes place. Because of the risky environment with permanent under- or overestimations, the most conservative of the rules, the regressive fundamental rule, gains strength and leads the spot rate back to the neighborhood of the fundamental value.

Moreover, each of the trading rules has a specific impact on the dynamics because of its lag structure. Depending on how the technical and fundamental trading rules are matched, different regimes of exchangerate fluctuations emerge. Since the agents in this model choose only from a set of six trading rules, the complexity of the exchange rate is inherently limited. An increase in the number of trading rules should augment the complexity of the dynamics.

According to (Goodhart, 1988), the empirical evidence indicates that the relationship between news and exchange-rate movements is rather fuzzy. Both systematic under- and overreaction to news are reported. For minor news, there is often no reaction at all. However, large price movements unrelated to any news are also apparent. Figure 6, showing the news arrival process (top) and the return dynamics (bottom), provides a first explanation for these findings: The amount and the frequency of extreme events in the simulated exchange-rate returns is much higher than in the incoming news. In general, one finds a relationship between the news arrival process and the returns. The traders anticipate this correctly, especially when a larger shock hits the market. Nevertheless, news is not incorporated immediately in the price. For minor shocks, the reaction can be quite contrary. As mentioned above, in period 4,130, the returns begin to fluctuate strongly without apparent reason. The explanation for these phenomena is rather simple: trading signals generated by technical trading rules can easily enforce or overcompensate the news effect.

Figure 6 also reveals that the model produces excess volatility. Guillaume et al. (1997) suggest to compute volatility as

$$v = (1/T) \sum_{t=1}^{T} |LogS_t - LogS_{t-1}| , \tag{16}$$

where T is the number of observations. Hence, the volatility is given as the average absolute logarithmic price change. In the same way, one can measure the volatility implied by the news process

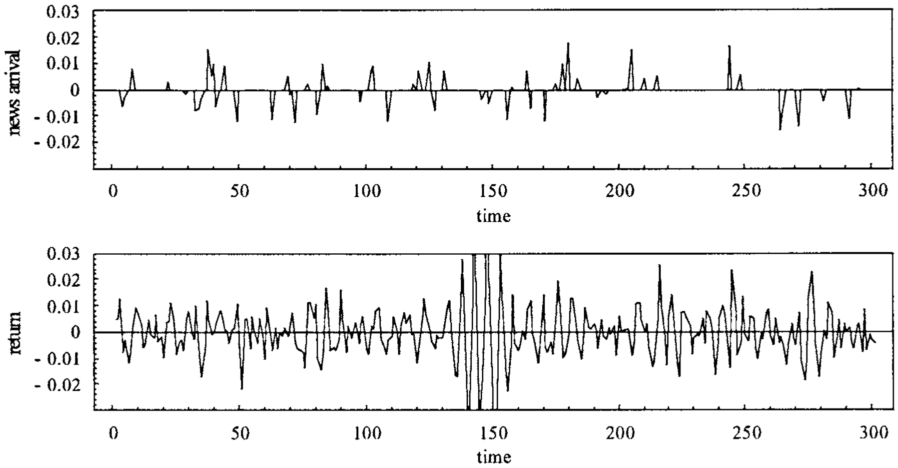

Figure 6. The relationship between news and returns for 300 periods. The top of the figure shows the news arrival process, the bottom the return dynamics, $T = 300$ observations, starting from $t = 4{,}000$.

$$\nu^F = (1/T) \sum_{t=1}^{T} |LogS_t^T - LogS_{t-1}^F| \, . \tag{17}$$

Excess volatility can then be defined as

$$e\nu = (\nu/\nu^F). \tag{18}$$

For our data set (5,000 observations), the volatility of the exchange rate is more than five times higher than the volatility of its fundamental ($e\nu = 5.18$).

A cointegration test checks if the exchange rate tracks its fundamental. The test consists of two steps (Engle and Granger, 1987). Cointegration requires that the variables to be integrated are of the same order. The first step is thus to determine the order of integration of the exchange rate and its fundamental. In our case, the ADF test should reveal that both variables are integrated of order 1 (unit roots!). Further, variables are cointegrated if there exists a linear combination that is stationary. Hence, the second step is to estimate the long-term equilibrium relationship between the variables. For this, one has to regress

$$S_t^F = aS_t + u_t \, , \tag{19}$$

and to apply the ADF test on the residuals

$$\Delta u_t^F = a_0 u_{t-1} + \sum_{j=1}^{l} a_j \Delta u_{t-j} \, , \tag{20}$$

If the residuals are stationary, one can conclude that the series are cointegrated.

Table II. Estimated kurtosis under time aggregation.

Time aggregation	1	5	10	25	50
Kurtosis	25.45	15.69	9.11	5.29	3.12

$T = 100{,}000$ observations, kurtosis for Normal distribution for all lags is 3.

We have used the following experiment to test for cointegration: The first 5,000 observations of the exchange rate and its fundamental have been split into 10 subsamples each of length 500. For each subsample, the exchange rate and its fundamental have been found to be integrated of the order 1. The critical value for stationary residuals at the 1% level is given by –3.73. Since the estimated t-values are in the range [–10.10, –8.54], we conclude that both time series are cointegrated.

Is the foreign-exchange market efficient or not? The answer to this question is not obvious. Our artificial market displays some mispricing, since the price fluctuations are only partially caused by the arrival of new information. Moreover, the interaction between the traders produces excess volatility (risk). Both phenomena typically have a negative impact on international trade. But in the long run, the exchange rate tracks its fundamental. The mispricing does not accumulate into lasting distortions.

3.3. RETURN DYNAMICS

A lot of empirical work describes the distribution of the returns. It is an important stylized fact that the distribution of the returns has fat tails (Guillaume et al., 1997). Relative to a Normal distribution, one finds a stronger concentration around the mean, more probability mass in the tails of the distribution, and thinner shoulders. Estimations of the kurtosis reveal the fat-tail property. Furthermore, the empirically observed kurtosis declines under time aggregation. Table II displays estimates of the kurtosis under time aggregation for 100,000 simulated exchange rates. The kurtosis of a Normal distribution is 3. Note that stronger outliers occur not only as a consequence of random shocks. If, for instance, a medium demand by chartists occurs with a low weight for fundamentalists, the price reaction is also strong.

An alternative way to identify fat tails is to determine the tail index. Empirical studies indicate that the distribution of large price changes roughly follows a power law (Guillaume et al., 1997; Lux and Marchesi, 1999; Farmer, 1999). Figure 7 compares the distribution of the returns and the scaling behavior for the simulated data with Normally distributed returns (with the same variance). The tail index α, given as $F(|return| > x) \approx cs^{-a}$, is estimated from the cumulative distribution of the positive and negative tails for normalized log-returns. The returns are normalized by dividing by the standard deviation. A regression on the largest 30% of the observations indicates a significant tail index of 3.38, which is in agreement with results obtained from the empirical data. The tail index for a Normal distribution, as

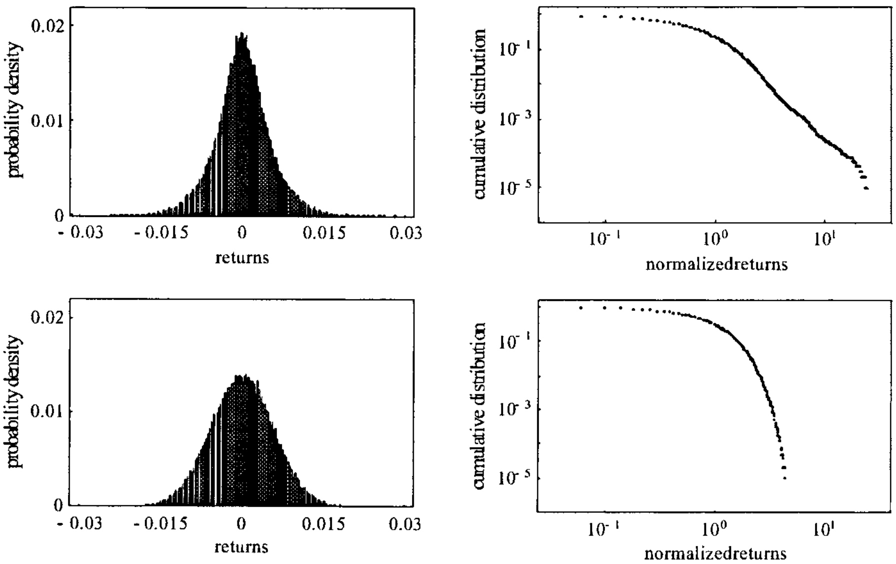

Figure 7. The distribution of the returns and the scaling behavior. The top left shows the distribution of the returns and the top right the scaling behavior of the cumulative distribution of the positive and negative tails for normalized log-returns, $T = 100{,}000$ observations, the bottom displays the same as before but for Normal distributed returns with identical variance.

seen for the slope in the right part of Figure 7, is clearly higher and less significant (the r-squared drops from 0.99 to 0.81).

Empirical results concerning serial autocorrelation of the returns of the exchange rates are not uniform. (Cutler et al., 1990) found that returns tend to be positively correlated at high frequencies and are weakly negatively correlated over longer horizons, thus exhibiting a mean reversion tendency. For other financial data, the mean reversion tendency is much stronger. The left hand of Figure 8 displays the autocorrelation function for the first 5,000 periods. The time series reveals significant mean reversion, since the autocorrelation for some lags lies clearly outside the 95% confidence intervals as given by $\pm 2/\sqrt{T}$, with T the number of observations, and the assumption that the returns are white noise. As already mentioned, the complexity of the exchange-rate movements is bounded because of the limited number of the trading rules. If one allows more rules, especially with a deeper lag structure, the autocorrelation declines.[5]

Since the work of Mandelbrot (1963), the high short-term autocorrelation of the volatility and its clustering in periods of high and low volatility are well known. Short-term autocorrelation is the result of market dynamics and not caused by a clustered arrival of news. That is, after new information has hit the market, the agents need some time to deal with the shock. Typically, technical traders jump on the bandwagon and reinforce the market movement. Moreover, technical trading rules are able to generate self-reinforcing trading signals of their own. The long-

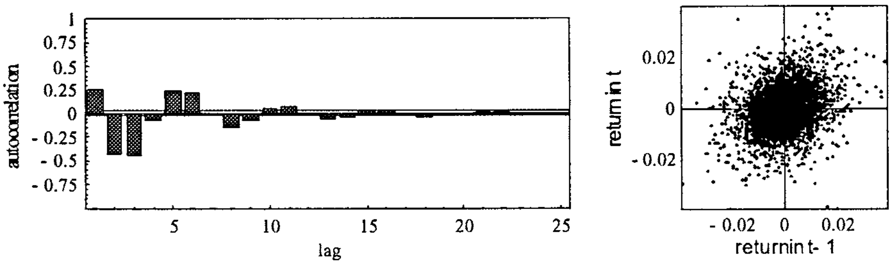

Figure 8. Autocorrelation function and phase space of returns. $T = 5{,}000$ observations, 95% confidence intervals are plotted as $\pm 2/\sqrt{T}$ (assumption of white noise).

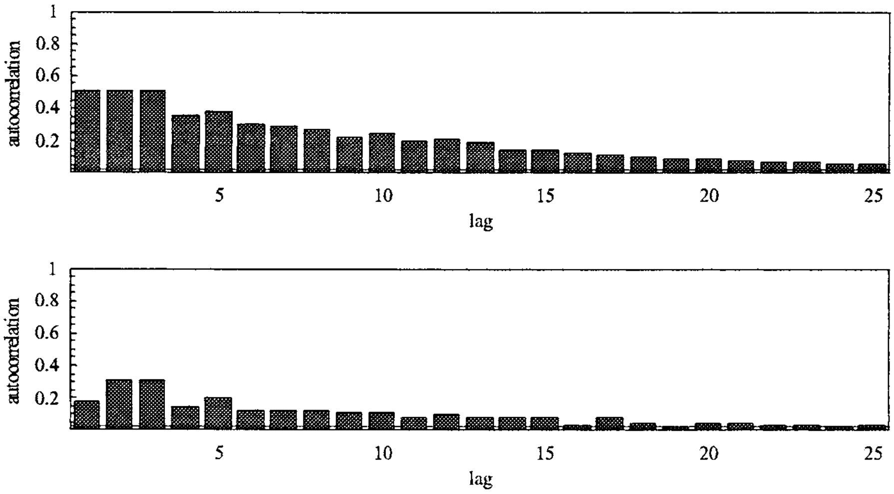

Figure 9. Autocorrelation function of trading volume and absolute returns. $T = 5{,}000$ observations, 95% confidence intervals are plotted as $\pm 2/\sqrt{T}$ (assumption of white noise).

run clustering may be explained by different degrees of uncertainty triggered, for example, by an oil-price shock or a political crisis. Figure 9 shows the autocorrelation function for the trading volume and the autocorrelation function for absolute returns. Both time series exhibit strong auto correlation. The short-term autocorrelation of the absolute returns is due to the trend-following behavior of the chartists; the long-term autocorrelation is a consequence of the variation in the matching of different trading rules, which results in periods of lower and higher volatility (compare also Figure 3). The volume of the trading positions of the chartists depends on their trading signals. Whenever the technical trading signals are clustered (the exchange rate movements are correlated), the trading volume is also clustered.

4. Conclusions

This paper presents a model where heterogeneous, boundedly rational market participants use a mix of technical and fundamental trading rules to determine their speculative investment positions. The trading rules and the selection of the strategies are not derived from a well-defined utility maximization problem but are based on empirical observations.

The model is able to replicate various well-known stylized facts. In its first moments, the computed time series looks random, as indicated by the unit roots. In the second moments, however, some (deterministic) patterns like volatility clustering or mean reversion are observable. The sources of the dynamics are endogenous processes, in particular the evolution of the agents' weighting scheme, rather than exogenous shocks.

Since we restrict our attention only to six prominent trading rules, there remains a higher structure in the simulated exchange-rate time series than typically observed in actual financial market data. This could be reduced by allowing for more heterogeneity. However, we abstain from this to see more clearly the major driving forces behind the foreign-exchange dynamics.

Notes

[1] For a brief introduction into technical analysis see Neely (1997).

[2] See Camerer (1995) or Slovic et al. (1988) for empirical surveys about such kind of expectation formation hypothesis.

[3] Note that if a low proportion of fundamentalists is confronted with a huge demand of chartists a very large price reaction is needed in order to match the demand. But this happens also from time to time in real financial markets. If the price reaction exceeds a certain threshold, trading is typically interrupted for a while to calm down the market.

[4] Since the model is calibrated to daily exchange rate fluctuations, this may be interpreted as a weekly performance evaluation.

[5] In addition, the technical trading rules used in this paper are to some degree similar. One reason for this is that we only allow for time lags of maximal 5 periods. However, in reality trading rules are more complex. Often, they are nonlinear and connected to threshold values (as for instance the head and shoulder pattern).

References

BIS (1999). *Central Bank Survey of Foreign Exchange and Derivatives Market Activity*. Basle.

Black, F. (1986). Noise. *Journal of Finance*, **41**, 529–543.

Brock, W. and Hommes, C. (1997). A rational route to randomness. *Econometrica*, **65**, 1059–1095.

Camerer, C. (1995). Individual decision making. In J. Kagel and A. Roth (eds.), *The Handbook of Experimental Economics*, 587–703.

Cutler, D., Poterba, J. and Summers, L. (1990). Speculative dynamics and the role of feedback traders. *American Economic Review*, **80**, 63–68.

Dawid, H. (1999). *Adaptive Learning by Genetic Algorithms*, 2nd edn. Springer, Berlin.

Dickey, D. and Fuller, W. (1981). Likelihood ratio test statistics for autoregressive time series with a unit root. *Econometrica*, **49**, 1057–1072.

Engle, R. and Granger, C. (1987). Co-integration and error correction: Representation, estimation, and testing. *Econometrica*, **55**, 251–276.

Farmer, D. (1999). Physicists attempt to scale the ivory towers of finance. *Computing in Science and Engineering*, **1**, 26–39.

Frankel, J. and Froot, K. (1986). Understanding the U.S. dollar in the eighties: The expectations of chartists and fundamentalists. *Economic Record*, **62**, 24–38.

Goodhart, C. (1988). The foreign exchange market: A random walk with a dragging anchor. *Economica*, **55**, 437–460.

Goodhart, C., McMahon, P. and Ngama, Y. (1993). Testing for unit roots with very high frequency spot exchange rate data. *Journal of Macroeconomics*, **15**, 423–438.

Grefenstette, J. (1986). Adaptive learning and expectational stability: An introduction. *IEEE Transactions on Systems, Man, and Cybernetics*, **16**, 122–128.

Guillaume, D., Dacorogna, M., Dave, R., Mueller, U., Olsen, R. and Picet, O. (1997). "From the bird's eye to the microscope: A survey of new stylized facts of the intra-daily foreign exchange markets. *Finance Stochastics*, **1**, 95–129.

Heiner, R. (1983). The origin of predictable behavior. *American Economic Review*, **73**, 560–595.

Kirman, A. (1993). Ants, rationality, and recruitment. *Quarterly Journal of Economics*, **108**, 137–156.

LeBaron, B. (2000). Agent based computational finance: Suggested readings and early research. *Journal of Economic Dynamics and Control*, **24**, 679–702.

Lux, T. and Marchesi, M. (1999). Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature*, **397**, 498–500.

Lux, T. and Marchesi, M. (2000). Volatility clustering in financial markets: A microsimulation of interacting agents. *International Journal of Theoretical and Applied Finance*, **2**, 675–702.

Mandelbrot, B. (1963). The variation of certain speculative prices. *Journal of Business*, **36**, 394–419.

Mandelbrot, B. (1997). *Fractals and Scaling in Finance*. Springer, New York.

Mitchell, M. (1996). *An Introduction to Genetic Algorithms*. MIT Press, Cambridge.

Murphy, J. (1996). *Technical Analysis of Financial Markets*. Institute of Finance, New York.

Neely, C. (1997). Technical analysis in the foreign exchange market: A layman's guide. *Federal Reserve Bank of St. Louis Review*, **79**, 23–38.

Slovic, P., Lichtenstein, S. and Fischhoff, B. (1988). Decision making. In R. Atkinson, R. Herrnstein, G. Lindzey and D. Luce (eds.), *Steven's Handbook of Experimental Psychology*, 673–738. Wiley, New York.

Taylor, M. and Allen, H. (1992). The use of technical analysis in the foreign exchange market. *Journal of International Money and Finance*, **11**, 304–314.