Inauguraldissertation
zur Erlangung des akademischen Grades eines
Doktors der Wirtschaftswissenschaften
des Fachbereichs Wirtschaftswissenschaften der
Universität Osnabrück

Chartists, Fundamentalists, and Exchange Rate Fluctuations

Frank H. Westerhoff

Dekan: Herr Prof. Dr. Dr. Giacomo Corneo

Referent: Herr Prof. Dr. Michael Braulke

Korreferent: Herr Prof. Dr. Lothar Knüppel

Tag der mündlichen Prüfung: 12.07.2002

Berichte aus der Volkswirtschaft

Frank Westerhoff

Chartists, Fundamentalists, and Exchange Rate Fluctuations

Shaker Verlag
Aachen 2002

Die Deutsche Bibliothek - CIP-Einheitsaufnahme

Westerhoff, Frank:
Chartists, Fundamentalists, and Exchange Rate Fluctuations /
Frank Westerhoff. Aachen : Shaker, 2002
(Berichte aus der Volkswirtschaft)
Zugl.: Osnabrück, Univ., Diss., 2002
ISBN 3-8322-0789-9

Printed in Germany.

ISBN 3-8322-0789-9
ISSN 0945-1048

Shaker Verlag GmbH • Postfach 101818 • 52018 Aachen
Telefon: 02407 / 95 96 - 0 • Telefax: 02407 / 95 96 - 9
Internet: www.shaker.de • eMail: info@shaker.de

Contents

1 Introduction

Exchange rate economics is in crisis. Frankel and Rose (1995) judge current macroeconomic approaches to exchange rates as empirical failures. Even leading economists in the field of international finance seem unable to benefit from their knowledge. For instance, Lyons (2001, page 1) puts his first "real-life experience" with the foreign exchange market in this way:

> "A friend of mine who trades spot foreign exchange for a large bank invited me to spend a few days at his side. That was ten years ago. At the time I considered myself an expert, having written my thesis on exchange rates. I thought I had a handle on how it worked. I thought wrong. As I sat there my friend traded furiously, all day long, racking up over $1 billion in trades each day ($US). This was a world where the standard trade was $10 million, and a $1 million trade was a "skinny one". Despite my belief that exchange rates depend on macroeconomics, only rarely was news of this type his primary concern. Most of the time he was reading tea leaves that were, at least to me, not so clear. The pace was furious – a quote every 5 or 10 seconds, a trade every minute or two, and continual decisions about what positions to hold. Needless to say, there was little time for chat. It was clear my understanding was incomplete when he looked over, in the midst of his fury, and asked me "What should I do?" I laughed. Nervously."

Structural exchange rate models seek to explain the behavior of exchange rates via macroeconomic variables. But - as the quotation of Lyons suggests - something may be missing in these models: the frantic trading of speculators! With a few exceptions, this aspect has not received much attention.

Already Keynes (1936) remarked that financial markets are ruled by the short-term speculative behavior of heterogeneous traders. Private investors pay close attention to non-fundamental information and are subject to waves of optimistic and pessimistic sentiments. Professional traders do not correct the vagaries of private investors. They are concerned with predicting asset prices under the influence of mass psychology for short-run profit making. Investments based on genuine long-term forecasts are rare since such an analysis is so difficult as to be scarcely practicable. Keynes concludes that the activity of speculators destabilizes financial markets and thus recommends the imposition of a substantial transfer tax to improve market efficiency.

Unfortunately, these arguments only obtained stronger support in the second half of the 1980s. The noise trader approach divides market participants into irrational and rational traders (Shleifer and Summers 1990). The demand of the first group is determined by popular models and responds to pseudo-signals, such as rumors. Since arbitrage positions of rational traders are limited, shifts in investor sentiments trigger price changes. A related area of research, based on empirical observations, assumes that agents rely on rather simple trading rules, such as technical and fundamental analysis, to determine their investment positions (Hommes 2001). Such models have the potential to mimic the behavior of financial prices, including exchange rates, quite well.

Though this work has several objectives, some of them deserve mention at the outset. One objective is to improve our understanding of exchange rate dynamics. For instance, we attempt to identify the driving forces of exchange rate fluctuations. Furthermore, we want to develop a model which generates "realistic" exchange rate dynamics. That is, the model should replicate stylized facts such as unit roots in the exchange rates, fat tails for returns, and volatility clustering. Another aim comprises the evaluation of certain means of control to reduce instability and avoid crashes.

Our study is organized as follows. Chapter 2 consists of two parts. We discuss the stylized facts of the foreign exchange market and survey papers which deal with the behavior of heterogeneous traders. The purpose of chapter 2 is to take the reader to the current state of research. Chapters 3 to 5 present our contributions to the literature. Each chapter is structured as a paper (including abstract, keywords and JEL classification) and can thus conveniently be studied separately. In brief: chapter 3 develops a nonlinear exchange rate model, in which traders have the choice between different kinds of trading rules to determine their investment positions. Even in a deterministic setting, the interaction between these rules produces complex dynamics. In chapter 4, the deterministic model is buffeted with shocks. Due to this extension, the model matches various stylized facts of exchange rate dynamics. Within this environment, we study the effectiveness of different kinds of market stabilization. Intervention strategies, such as "leaning against the wind", may reduce exchange rate variability as long as the trading behavior of the agents is correlated. Clearly, this analysis is preliminary. But one should note that this is one of the very first attempts to deal with such an issue. Chapter 5 adds to the literature in that psychological factors influencing the decision-making process of the traders are modeled explicitly. As a result, we are able to simulate artificial exchange rate data which is difficult to distinguish from real exchange rate data. Finally, chapter 6 offers some conclusions.

2 The Foreign Exchange Market

It is well known that empirical exchange rate models based on fundamentals perform poorly (Meese and Rogoff 1983). This has led to an ongoing search for better approaches to modeling exchange rate dynamics. Recently some studies have provided suggestive evidence that the act of trading, apart from any news regarding fundamentals, may be an important source of exchange rate movements.

In this chapter, we survey three related fields of this kind of research: the noise trader approach, chartists-fundamentalists models and agent-based computation. These contributions are based on the observations that at least some agents rely on rather simple trading strategies to determine their investment positions. The interaction between the traders is seen as a key element to explain the dynamic behavior of exchange rates. Chapter 2 is structured as follows. Section 2.1 describes the foreign exchange market. We characterize how the foreign exchange market operates and discuss some statistical properties of exchange rate fluctuations. Section 2.2 surveys the above literature and section 2.3 gives a short look at our own work.

2.1 Description of the Foreign Exchange Market

The foreign exchange market is the broadest, most active financial market in the world. It operates 24 hours, stopping only for weekends. The first transactions of the week are settled at 23:00 Greenwich Mean Time (GMT) on Sunday with the opening of the Asian markets. The trading stops at about 23:00 GMT on Friday when the US West Coast markets close. The foreign exchange market is structured as an electronic/telephone market. The prices that the main players are prepared to bid and offer can be observed by watching the relevant screen pages of financial news agencies such as Reuters. Given such publicly available information on the bid-ask rates on offer, deals can be either made over the telephone or via automated dealing systems. The amount and the price of each transaction remains the private information of the two parties concerned.[1]

[1] More information about the market microstructure is, for instance, provided by Lyons (2001), Frankel et al. (1996), and Flood (1991).

Section 2.1 aims at describing the foreign exchange market by reviewing some of its most important stylized facts. At this point, the nature of the analysis is just descriptive. Explanations for these phenomena are provided later. The remainder of this section is organized as follows. Section 2.1.1 presents our data set, section 2.1.2 summarizes the stylized facts and section 2.1.3 deals with technical analysis.

2.1.1 The Data Set

Since the late 1980s, intra-daily data of the foreign exchange market has been accessible. Relative to daily data, the number of data points available intra-daily is 100 to 1,000 times larger. Guillaume et al. (1997) report that on the DEM/USD market 4,500 quotes per day are recorded. That is, there is an average of 3 to 4 new quotes per minute. This average can rise to 15-20 quotes per minute during the busiest periods. Guillaume et al. focus on the period 01/01/1987-12/31/1993, with a total of 8,238,532 observations for the DEM/USD rate. Such a data base naturally opens the door to new insights. For example, the impact of news on the exchange rate can almost be analyzed in continuous time. Unfortunately, such time series are very expensive. Moreover, data on foreign exchange transactions is not made public (Cai et al. 2001). Historically, this was due to the bilateral nature of direct dealing trades where only the two participants to the trade knew the quantities traded. More recently with the advent of electronic trading, it is feasible to provide data representing a large cross-section of the market, but the firms providing the major dealing systems such as Reuters do not make such data public in order to preserve the confidentiality of the trading parties.

The Deutsche Bundesbank kindly provided us with a DEM/USD and DEM/JPY exchange rate time series. The data set under consideration extends from 01/01/1974-12/31/1998. The beginning of the period coincides with the breakdown of the Bretton Woods System. The end of the period is marked by the launch of the Euro. Each time series contains 6,264 daily exchange rate observations (5 entries per week). The exchange rates are given with 5 digits of precision (for example 2.7245). Further information regarding the data sampling can be found in the Devisenkursstatistik of the Deutsche Bundesbank (1998). Figure 1 shows the DEM/USD and DEM/JPY exchange rates, respectively.

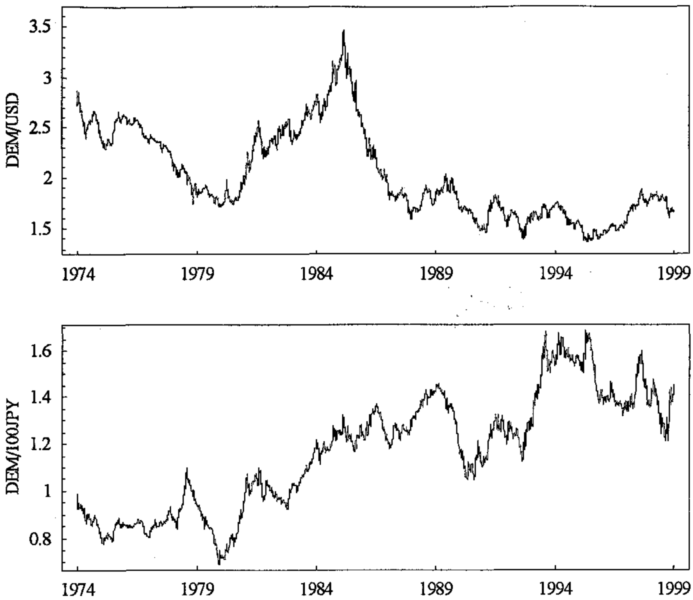

Figure 1: The Data Set. Exchange rates over the period 01/01/1974-12/31/98, 6,264 daily observations, source: Deutsche Bundesbank (1998).

2.1.2 Some Stylized Facts

This section presents a couple of stylized facts which describe the foreign exchange market with respect to our work. For a better understanding, we try to replicate some of these findings with our data set. For this, we also repeat the relevant methods of time series analysis. Note that it is not our intention to question the stylized facts nor to improve any testing procedure. These observations have been independently confirmed in many empirical studies. We just want to clarify what is meant by them and how the tests are carried out.

General surveys on the foreign exchange market are provided by Andersen and Bollerslev (1998), Goodhart and O'Hara (1997), Guillaume et al. (1997), Pagan (1996), de Vries (1994), and Goodhart and Giugale (1993). Useful introductions into time series analysis can be found in Campbell et al. (1997), Enders (1995), and Hamilton (1994).

2.1.2.1 Unit Roots

One stylized fact of the empirical literature is that exchange rate time series display unit roots (Goodhart et al. 1993a, Goodhart and Taylor 1992). The augmented Dickey-Fuller (ADF) unit root testing procedure (Dickey and Fuller 1979, 1981) tests the null hypothesis that any shock to the exchange rate is permanent against the alternative hypothesis that a shock is only temporary. In particular, one estimates the regression

$$\Delta S_t = a_0 S_{t-1} + a_1 \Delta S_{t-1} + a_2 \Delta S_{t-2} + \ldots + a_n \Delta S_{t-n}, \quad (1)$$

where S denotes the exchange rate and ΔS the first difference. The lag length n can be optimized with respect to desirable statistical properties. The coefficients a_0-a_n are estimated via OLS regression. The null hypothesis of a unit root is rejected, if a_0 is negative and significantly different from zero.

Table 1 shows the estimation results for a four-lag specification. The theoretical t-values for a_0 are -1.62, -1.95 and -2.58 at the 10, 5 and 1 percent significance levels, respectively. Since the empirical t-values clearly exceed the theoretical t-values, it is not possible to reject the hypothesis of the existence of a unit root. This result is often interpreted as evidence for a random walk behavior of the exchange rate.

	a_0	a_1	a_2	a_3	a_4
DEM/USD	-0.000126 (-1.432)	-0.031369 (-2.484)	-0.005145 (-0.407)	0.034262 (2.716)	-0.003293 (-0.261)
DEM/JPY	0.000046 (0.577)	0.015582 (1.232)	0.015704 (1.243)	-0.021596 (-1.710)	-0.007997 (-0.633)

Table 1: Unit Root Test. Data set of figure 1, t-values are given in parenthesis.

2.1.2.2 Excess Volatility

Whether exchange rates are too volatile with respect to the behavior of their underlying determinants is one of the most important issues in the international finance debate. The issue is crucial for policy purposes: if exchange rates are too volatile with respect to a reasonable benchmark in an efficient market, then there may be grounds for "throwing sand in the wheels" of currency markets, so as to slow down changes in exchanges rates and keep them in line with those of their fundamentals (Tobin 1978, Eichengreen et al. 1995).

Volatility measures are mainly based on the return process. Returns are defined as the change of Log prices

$$r_t = LogS_t - LogS_{t-1}. \tag{2}$$

Guillaume et al. (1997) suggest computing the volatility as

$$v = (1/T)\sum_{t=1}^{T} | r_t |, \tag{3}$$

where T is the number of observations. The volatility is given as the average absolute return over the sample period. In our case, the volatility for the DEM/USD time series is calculated as v=0.0050 and for the DEM/JPY time series as v=0.0044. Alternatively, the volatility may be measured by the sample variance (compare the next section).

Cai et al. (2001), Bartolini and Giorgianni (2000), and Andersen and Bollerslev (1998) find broad evidence of excess volatility for major currencies over the post Bretton Woods area. For them, such volatility estimates are too high to be justified by fundamental shocks. Although this seems to be the majority opinion, some studies are opposed to the notion of excess volatility. Eichenbaum and Evans (1995) identify monetary shocks as an important source of exchange rate fluctuations. According to them, such relationships are often underestimated.

2.1.2.3 Fat Tails

Relative to the distribution of the exchange rates, the distribution of the returns is more symmetric. In addition, the process of the returns is clearly stationary. Hence, the latter is usually preferred to describe foreign exchange dynamics. Figure 2 displays the return dynamics for our data set in the time domain. The top contains DEM/USD returns and the bottom DEM/JPY returns.

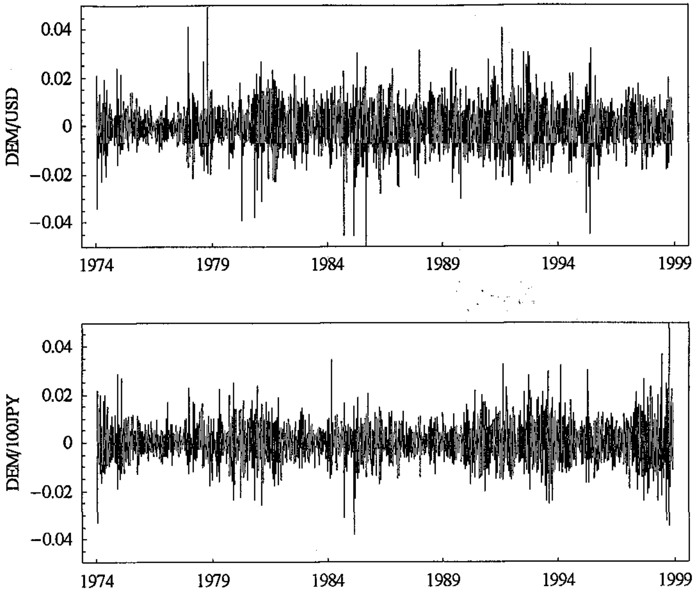

Figure 2: The Returns. Data set of figure 1, 6,263 daily observations.

In the first step, the distribution of the returns is typically characterized by its first four moments. The sample mean is defined as

$$\bar{r} = \frac{1}{T} \sum_{t=1}^{T} r_t , \tag{4}$$

the sample variance as

$$\sigma^2 = \frac{1}{T} \sum_{t=1}^{T} (r_t - \bar{r})^2 , \tag{5}$$

the sample skewness as

$$SK = \frac{1}{T\sigma^3} \sum_{t=1}^{T} (r_t - \bar{r})^3 , \tag{6}$$

and the kurtosis as

$$K = \frac{1}{T\sigma^4} \sum_{t=1}^{T} (r_t - \bar{r})^4 . \tag{7}$$

Table 2 reports estimates of the moments together with the minimum and maximum return values for our data set. In comparison, the normal distribution has skewness equal to 0 and kurtosis equal to 3. Fat-tailed distributions with extra probability mass in the tail area have higher or even infinite kurtosis. Excess kurtosis is thus defined to be sample kurtosis less 3. Our calculations reveal that both time series exhibit signs of fat-tailedness. Moreover, largest returns can exceed the 5 percent level.[2] Guillaume et al. (1997) report similar values for other currencies.

	min	mean	max	variance	skewness	kurtosis
DEM/USD	-0.0575	-0.0001	0.0495	0.000048	-0.1579	7.043
DEM/JPY	-0.0384	0.0001	0.0894	0.000038	0.7601	13.982

Table 2: Some Descriptive Return Statistics. Data set of figure 1.

In the early literature, fat tails have been identified with excess kurtosis. However, for financial data the kurtosis is a very unreliable estimator of fat-tailedness. Fortunately, recent literature (Reiss and Thomas 2001) provides a better description of the behavior in the distribution's extreme parts. Since the decline of probability mass in the outer parts roughly follows a power law, the distribution in the tails can be approximated by

$$F(|r|>x) \approx cx^{-a}, \qquad (8)$$

where a is the so called tail index. A lower a indicates fatter tails. The tail index 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. Farmer (1999) and Lux and Ausloos (2000) report that a usually ranges between 2 and 4 for major currencies. Reiss and Thomas (2001) show that the moments of a distribution higher than its tail index are not bounded. Therefore, the tail index should be preferred to the kurtosis to identify fat tails.

Figure 3 compares the distributions of the returns and their scaling behavior of our data set with normally distributed returns. The shape of the distribution appears well behaved: the histograms of the returns are uni-modal, almost symmetric and bell shaped. A regression on the largest 30 percent of the observations delivers a significant tail index of 3.68 for the DEM/USD and of 3.39 for the DEM/JPY time series. The tail index of the Normal distribution is clearly higher (compare the slope in figure 3).

[2] On 10/07/1998, the yen price of a dollar fell from about 134 to 120 (10.8 percent). On the next day, the yen declined further to 111 (Cai et al. 2001). Such extreme events rarely occur. The future may even show stronger exchange rate movements!

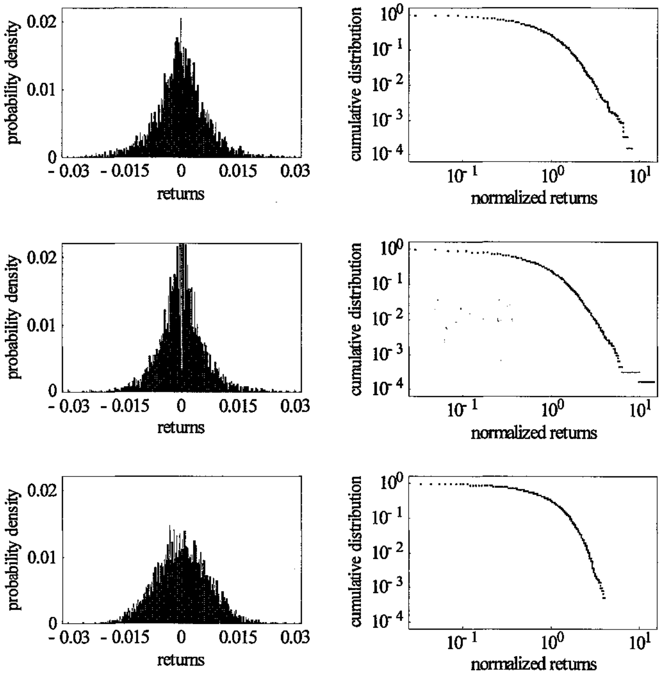

Figure 3: Distribution of Returns and Scaling Behavior. Data set of figure 1, top (middle, bottom) DEM/USD (DEM/JPY, Normal distribution) data.

Hill (1975) provides a tail index estimator which is easy to compute. The sample elements are put in descending order: $X_T > X_{T-1} > ... > X_{T-k} > ... > X_1$, with k being the number of observations located in the tail. The Hill tail index estimator is derived as

$$a^H = (\frac{1}{k} \sum_{i=1}^{k} Log\, X_{T-i+1} - Log\, X_{T-k})^{-1}. \quad (9)$$

The estimate depends on an appropriate choice of the tail region (the cut-off value k). Lux (2001) compares different techniques to determine optimal cut-off values. Methods for endogenous selection of k basically tend to choose estimates around "stable" regions, i.e. near a plateau.

Since there is no indication of systematic differences between positive and negative price changes, both tail parts are merged by using absolute returns. Figure 4 displays Hill tail index estimates for tail fractions between 0 and 10 percent. Frequently used tail fractions include 2.5 and 5 percent. Note that for the DEM/USD data, a plateau appears between 5 and 7 percent and for DEM/JPY data between 2.5 and 5 percent, indicating a tail index around 3.6. The results are thus close to the previous estimates. To sum up, the tail index seems to range between 3 and 4, implying the existence of the first and second moment, but non-convergence of the fourth moment.

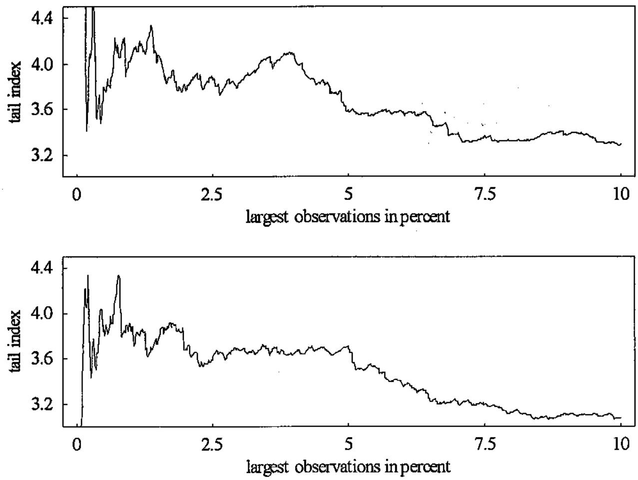

Figure 4: Hill Tail Index Estimator. Data set of figure 1, top DEM/USD and bottom DEM/JPY.

2.1.2.4 Mean Reversion

According to Cutler et al. (1990, 1991) exchange rate returns tend to be positively serially correlated at high frequency and weakly negatively serially correlated over long horizons. This is called mean reversion. Mean reversion in financial markets is reported for quite different time horizons. The studies of Cutler et al. explore monthly data, whereas Mantegna and Stanley (2000) find significant correlation over a couple of minutes for intra-daily data, and Campbell et al. (1997) over one or two lags for data on a daily level. Clearly, mean reversion implies predictability. At first sight, this seems to be inconsistent with efficient markets. But as noted by Campbell et al. a certain degree of predictability may be necessary to reward investors for bearing certain dynamic risks.

Serial correlation measures the correlation between two observations of the same time series at different dates. The sample autocorrelation function for the returns is defined as

$$ACF(k) = \frac{(1/T)\sum_{t=1}^{T-k}(r_t - \bar{r})(r_{t+k} - \bar{r})}{(1/T)\sum_{t=1}^{T}(r_t - \bar{r})(r_t - \bar{r})}, \qquad (10)$$

where k indicates the lag size. If the data is generated by Gaussian white noise, then the ACF should lie between $\pm 2/\sqrt{T}$ about 95 percent of the time.

Figure 5 shows the autocorrelation function for DEM/USD returns in the top and the autocorrelation function for DEM/JPY returns in the bottom. Overall, evidence for mean reversion within our data set is very weak. For almost all lags the autocorrelation is not significant. The first lag is positive as suggested by the mean reversion hypothesis only in the DEM/JPY case. In the DEM/USD market, the autocorrelation for k=1 is negative. A daily data sampling may not be frequent enough to detect mean reversion.

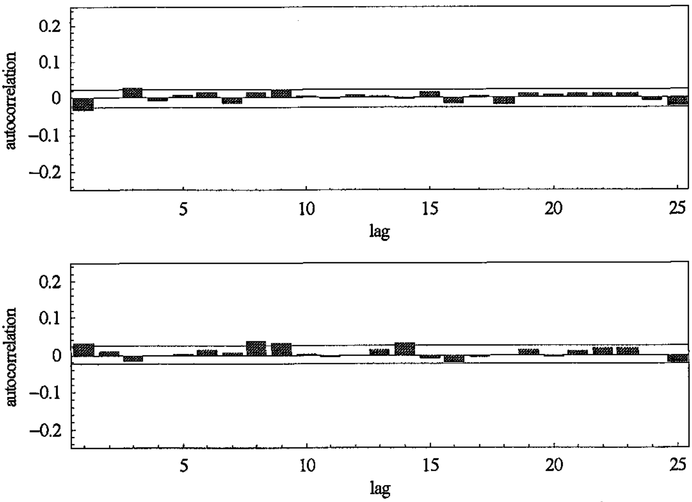

Figure 5: Mean Reversion. Data set of figure 1, 95 percent confidence bands as $\pm 2/\sqrt{T}$.

2.1.2.5 Volatility Clustering

One of the most robust stylized facts observed in the foreign exchange market is volatility clustering (Mandelbrot 1963). Visual inspection of figure 2 already reveals that periods of low volatility alternate with periods of high volatility. There are several suitable ways to quantify volatility clustering. For instance, autocorrelation functions can be derived by using absolute or squared returns as volatility measures. For both measures, volatility is persistent (figure 6). For absolute returns, the autocorrelation is around 0.2 in the short run and then slowly decays. The autocorrelation of squared returns is somewhat lower but still significant over long horizons.

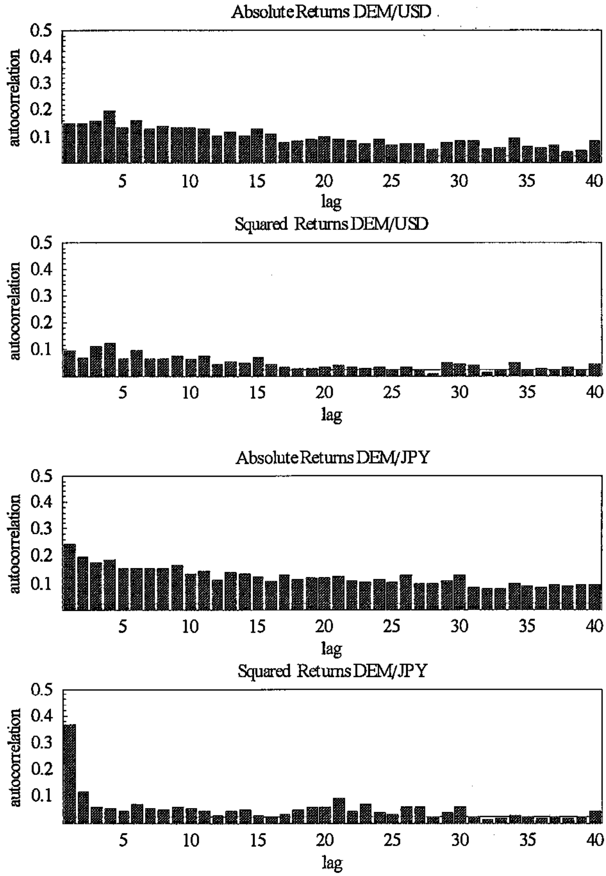

Figure 6: Volatility Clustering. Data set of figure 1, 95 percent confidence bands as $\pm 2/\sqrt{T}$.

Note that volatility persistence is significant over an extended time horizon with autocorrelation coefficients of absolute returns not even decaying to zero when considering 250 lags. Such long term dependencies can be quantified by calculating Hurst exponents (Hurst 1951, Mandelbrot 1971). Let

$$\overline{X}_N = (1/N)\sum_{t=1}^{N} X_t \tag{11}$$

and

$$S_N = \sqrt{(1/N)\sum_{t=1}^{N}(X_t - \overline{X}_N)^2} \tag{12}$$

be the mean and the standard deviation of a time series until date *N*. The accumulated deviations of the time series from its mean can be expressed as

$$Y(t,N) = \sum_{t=1}^{N}(X_t - \overline{X}_N). \tag{13}$$

Now, a so called re-scaled range statistic is obtained as

$$\frac{R}{S}(N) = \frac{\max\limits_{0\le t\le N} Y(t,N) - \min\limits_{0\le t\le N} Y(t,N)}{S(N)}. \tag{14}$$

Thus, the *R/S* statistic is the range of partial sums of deviations of a time series from its mean, re-scaled by its standard deviation. Finally, the Hurst exponent is defined as

$$\frac{R}{S}(N) = (\frac{N}{2})^H. \tag{15}$$

The Hurst exponent is estimated from the slope of *R/S* plotted against *N*/2 on a Log-Log scale. Hurst exponents larger than 0.5 indicate persistence, whereas Hurst exponents smaller than 0.5 indicate anti-persistence. H=0.5 hints at Brownian motion.

Figure 7 shows the determination of Hurst exponents for DEM/USD and DEM/JPY absolute returns. Unfortunately, the algorithm for computing Hurst exponents operates very slowly for large data sets. Thus, the first 6,000 observations have been truncated into 6 subsamples, each of length 1,000. Regressions deliver Hurst exponents for the DEM/USD data as H=0.72 and for the DEM/JPY data as H=0.70. Similar values are reported for other financial markets (Lux and Ausloos 2000).

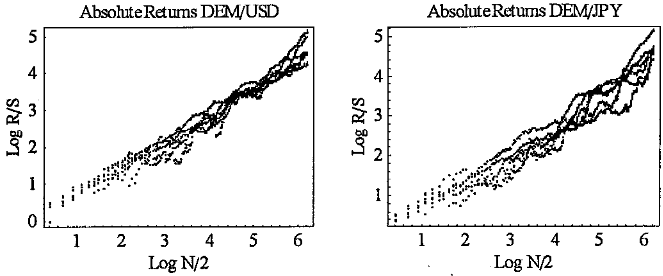

Figure 7: Hurst Exponents. Data set of figure 1, the first 6,000 observations are subdivided into 6 blocks, each of size 1,000.

The degree of persistence appears to be higher in absolute returns than in squared returns (compare figure 6). More specifically, for increasing powers to absolute returns, the memory effect declines. Such variations in the scaling exponent for various powers is called multi-scaling (Lux and Ausloos 2000). Figure 8 shows Hurst exponents for DEM/USD and DEM/JPY returns with powers ranging from 1 to 10.

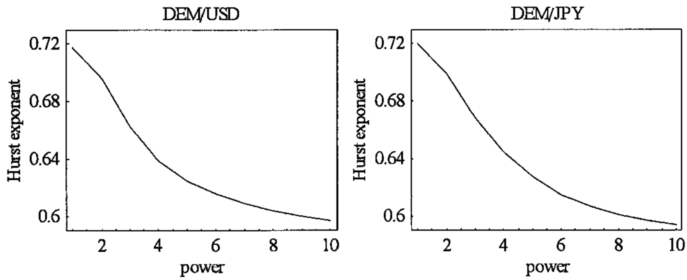

Figure 8: Multi-Scaling. Data set of figure 1, the first 6,000 observations are subdivided into 10 blocks, each of size 600.

The clustering of volatility in periods of high and low volatility is also captured by the famous GARCH class of models. ARCH models (Autoregressive Conditional Heteroskedasticity) were introduced by Engle (1982) and generalized as GARCH models (Generalized ARCH) by Bollerslev (1986). These models are widely used in financial time series analysis. For instance, the popular GARCH (1,1) specification is given as

$$r_t = c_1 + \varepsilon_t, \tag{16}$$

$$\sigma_t^2 = c_2 + \alpha \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2. \tag{17}$$

The first equation gives the conditional mean as a function of a constant c_1 and an error term ε. In the second equation, the conditional variance is a function of a constant c_2, news about the volatility from the previous period (the ARCH term), and the last period's forecast variance (the GARCH term).

According to (17), this period's variance is predicted by forming a weighted average of a constant, the volatility observed in the last period and the forecasted variance from the last period. The coefficients of the model are estimated by the method of maximum likelihood with the help of iterative algorithms. Table 3 displays our estimation results for the variance equation. It is a common finding that the sum of the ARCH and GARCH coefficients is close to one, implying – again – that volatility shocks are quite persistent (similar values are reported by Goodhart and O'Hara 1997, and Andersen and Bollerslev 1997).

	Constant	ARCH	GARCH
DEM/USD	0.00000060	0.112	0.883
DEM/JPY	0.00000057	0.114	0.876

Table 3: GARCH Estimates. Data set of figure 1.

2.1.2.6 Bubbles

Probably the most popular example of a bubble is the exchange rate behavior of the US dollar in the 1980s. Over the period 1981-1984, when real interest rates in the United States rose above those among trading partners, the dollar appreciated from DM 1.93 to DM 2.67. At times, however, the path of the dollar has departed from what would be expected on the basis of macroeconomic fundamentals. The most dramatic episode is the period from June 1984 to February 1985. The dollar appreciated further from DM 2.67 to DM 3.46 over this interval, even though the real interest differential had already begun to fall. Other observable factors, suggested in standard macroeconomic models, at that time were also moving in the wrong direction to explain the dollar rise. Taking these observations into account, Frankel and Froot (1986, 1990) argue that at least 20 percent of the appreciation are due to non-fundamental forces. After February

1985, the bubble burst. By the end of 1987, the exchange rate dropped below DM 1.60 (compare also figure 1).

Cointegration analysis is, in principle, a tool to detect such long term bubbles. Variables are cointegrated if there exists a linear combination that is stationary (Engle and Granger 1987). If the exchange rate tracks its fundamental, then the difference between the exchange rate and its fundamental should be stationary. In a broader sense, an equilibrium exists if the difference between the variables does not become too large.

To test for cointegration, Enders (1994) suggests the following procedure. Cointegration requires that the variables to be integrated are of the same order. Hence, the first step is to determine the order of integration. This can be done with the augmented Dickey-Fuller (ADF) test. If both variables are integrated of the order 1 (i.e. they display unit roots), the next step is to estimate the long-run equilibrium relationship between them. Let S be the exchange rate and S^F its fundamental. Then, regress

$$S_t^F = aS_t + u_t, \tag{18}$$

and apply the ADF test on the residuals u. Using, for instance, a four lag specification without drift delivers

$$\Delta u_t = a_0 u_{t-1} + a_1 \Delta u_{t-1} + a_2 \Delta u_{t-2} + a_3 \Delta u_{t-3} + a_4 \Delta u_{t-4}. \tag{19}$$

If the residuals are stationary (a_0=0 can be rejected), one can conclude that both time series are cointegrated. In this case, the exchange rate tracks its fundamental.

Surveying the empirical literature, Taylor (1995) states that evidence for cointegration between the exchange rate and its long-run equilibrium, derived from structural models such as the monetary approach, seems to be weak. While MacDonald (1999) finds some evidence for cointegration, this is strongly rejected by Groen (1999). These divergent results may be explained by a finding of Obstfeld and Taylor (1997). They present evidence that price differentials net of transactions costs are in general substantially eliminated in months rather than in years – but that costs of international trade result in "bands of inaction" within which relative international prices can fluctuate with no central tendency. Besides such lasting deviations, bubbles are also observable in the short run, e.g. as an overreaction to news. This is discussed in the next section.

2.1.2.7 The Impact of News

What is news? According to Guillaume et al. (1997), news is a very broad concept covering events like a phone call of a customer who wants to make a large transaction, a conversation between foreign exchange dealers, forecasts of technical traders or research centers, or general economic and political news releases. News is obviously very difficult to quantify.

Nevertheless, the impact of news on exchange rates is a heavily debated topic. Early contributions often reject a clear relationship between news announcements and exchange rate changes: even the largest daily price changes generally could not be associated with any significant news. For instance, Goodhart (1988) finds no significant response of spot rates to unanticipated changes in interest rates. He states that in the short run, the empirical evidence indicates that the reaction of exchange rates to news is an underreaction rather than an overreaction.

Recent papers focussing on high frequency data deliver a slightly different picture. Goodhart and Figliuoli (1991) have been among the first who relate movements in the exchange rate at high frequency to Reuters news announcements. These news messages appear in real-time on the screen of professional traders. Analyzing 3 days in 1987, it proved to be extremely difficult to relate movements in the foreign exchange market to such announcements, either prior, concurrently or immediately afterwards.

Similarly, Goodhart et al. (1993b) explore high frequency data extending over a period of 8 weeks. They detect an influence of news on both the level and the volatility of exchange rates. But these effects are not permanent. Both the amount of volatility and the level of the exchange rate tend to return to their pre-announcement values. This is also confirmed by Alameda et al. (1998). Andersen and Bollerslev (1998) find a somewhat stronger support. Largest returns appear to be linked to the release of public information such as certain macroeconomic announcements. More specifically, information arrivals induce abrupt price changes. Although the average price movement is typically attained within minutes, volatility remains elevated for several hours. However, for Andersen and Bollerslev these effects are secondary when explaining overall volatility. Finally, Guillaume et al. (1997) observe that distinct and relatively large price changes unrelated to any news are also apparent. They venture that the price formation process is at least partially independent of the presence or absence of news.

Empirical attempts to link exchange rate movements to news effects are hampered by the difficulty in extracting the unexpected component of the news. For this reason, Cheung and Chinn (2001) try to derive the impact of news from a survey study conducted among US-based foreign exchange traders. Their findings are as follows. News about macroeconomic variables is rapidly incorporated into exchange rates. For the majority of the respondents, the bulk of the adjustment takes place within one minute. About 1/3 of the traders even indicate that full adjustment takes place in less than 10 seconds. In these cases, even minute by minute data might not catch this news effect. The importance of individual macroeconomic variables shifts strongly over time. For instance, unemployment rates are regarded as important in certain times, but in other times they obtain only minor attention. One exception are interest rates, which are always seen as important.

2.1.2.8 Speculative Trading

The foreign exchange market is the largest financial market in the world. In its triennial central bank survey of foreign exchange and derivatives market activity, the Bank for International Settlement (BIS 1999) reports that the global turnover in traditional foreign exchange market segments reached an estimated daily average of US $1.5 trillion in April 1998. The market is strongly expanding. At April 1998 exchange rates, transactions rose about 46 percent in the three year period beginning in April 1995, after expansions of 29 percent between 1992-1995 and 33 percent between 1989-1992.

The US dollar is by far the most actively traded currency, followed by the Deutsche mark, the Japanese yen and Pound sterling. Together, these four currencies comprise 3/4 of total transactions. A geographical breakdown reveals a strong concentration of foreign exchange market turnover in a few centers. The five largest centers, the United Kingdom (32 percent), the United States (18 percent), Japan (8 percent), Singapore (7 percent) and Germany (5 percent), account for 69 percent of total transactions. In addition, there exists a high concentration in individual market centers. In the United

Kingdom and in the United States the market share of the top 10 dealers is around 50 percent. In medium sized markets, concentrations tend to be even higher. In France, the top 10 institutions alone account for 80 percent of the trading volume.

The foreign exchange market seems to be highly speculative. First, the overwhelming part of the trading volume reflects very short-term transactions. For example, operations of intra-day traders account for 75 percent of the market volume. Second, only 15 percent of the trading volume is due to none-financial customers. Third, international trade transactions represent merely 1 percent of the total volume. Schulmeister (1988) estimates that speculative transactions are roughly 20 times higher than the sum of world trade and international portfolio investment.

Unfortunately, the only available time series providing data on trading volumes are rather short and not representative (Goodhart and O'Hara 1997). However, financial markets often display strikingly similar stylized facts. Exploring daily stock data (both individual stocks and an index), Brock and LeBaron (1996) find positive autocorrelation functions of trading volumes with slowly decaying tails. In addition, using absolute or squared returns as a volatility measure, the cross-correlation is nearly zero for volatility with past and future volumes and positive for volatility with current volumes.

2.1.2.9 Trading Rules

Survey studies like Menkhoff (1997) or Taylor and Allen (1992) provide information about the behavior of speculators. Surprisingly, the traders rely on rather simple trading rules when determining their speculative investment positions. The trading rules belong either to the class of fundamental or technical analysis. Both concepts appear to be equally important. Fundamental analysis is built on the premise that the exchange rate converges towards its fundamental value. Fundamental trading rules suggest buying (sell) foreign currency when the exchange rate is below (above) its fundamental value. From an economist's point of view, such a behavior appears to be rational.

Rules based on technical analysis derive trading signals out of past price movements. For instance, an increase in the exchange rate may reflect a temporary trend so that a buy signal is triggered. Although such technical concepts are widely used in financial markets, they are not very familiar to the economic profession. For this reason, technical analysis is discussed more extensively in section 2.1.3.

2.1.2.10 Chaos and Nonlinearities

Simply speaking, a nonlinear dynamical system is a system whose time evolution equations are nonlinear. Suppose R is the response of a system to a stimulus S. If one doubles the stimulus, a linear system will have the response 2R. For a nonlinear system, however, the response will be larger or smaller than 2R. Nonlinearity is a necessary condition for chaos, but not a sufficient one.

Chaos is typically associated with three properties (for a short introduction see Baumol and Benhabib 1989, a deeper discussion follows in chapter 3). First, the trajectory of a deterministic process appears highly irregular. Second, the trajectory reacts in a sensitive manner to a microscopic change in the initial conditions. Third, the trajectory displays some long-term order in phase space (emergence of an attractor). Sensitivity to initial conditions and the order in phase space can be quantified by estimating the largest Lyapunov exponent and the correlation dimension, respectively. The correlation dimension also indicates the number of variables needed to reconstruct the dynamics. If experimental data shows a low dimensional attractor, relatively simple models may be sufficient to replicate the dynamics.

Interest in nonlinear chaotic systems has experienced a tremendous rate of development in the recent past (Barnett and Serletis 2000). One reason for this interest is the ability of such processes to generate output that mimics the output of stochastic systems, thereby offering an alternative explanation for the behavior of price dynamics. In particular, if the foreign exchange market possesses a low dimensional attractor, then the dynamics are mainly driven by an endogenous system with some small stochastic perturbations. Rather than the news process itself, it is the complex nonlinear interaction and learning process of traders that would be responsible for the large and unpredictable movements of the exchange rates.

Barnett and Serletis (2000) and Brock and Potter (1993) review some tests for nonlinearity and chaos. On the basis of recent empirical contributions both studies conclude that there is clear evidence of nonlinear dependence and some evidence of chaos in financial markets. Studies in favor of chaos include Bajo-Rubio et al. (1995), Scheinkman and LeBaron (1989), and Frank and Stengos (1989). A more pessimistic view is held by Guillaume (1995), who is not able to detect a low dimensional attractor.

However, one should note that the presence of dynamic noise makes it difficult and perhaps impossible to distinguish between (noisy) high dimensional chaos and pure randomness. The estimates of the correlation dimension and Lyapunov exponents of an underlying unknown dynamical system are sensitive to dynamic noise, and the problem grows as the dimension of the chaos increases. Clearly, low dimensional systems buffeted with random perturbations may be identified as stochastic systems. Moreover, the analysis is often further hampered by a low sample size of available data (Barnett and Serletis 2000).

2.1.2.11 Summary

The high trading volume of the foreign exchange market is strongly influenced by speculative activity. Surprisingly, the investment decisions of the agents are based on rather simple trading strategies. The interaction between the traders together with the news arrival process generates complex dynamics. On first sight, the exchange rate path appears as a random walk: exchange rates display unit roots and raw returns only show slight deviations from pure randomness. But on second sight, one can detect some finer structure in the time series such as volatility clustering. In addition, some deviations from IID-ness are indicated by the methods of nonlinear time series analysis.

Overall, the exchange rate volatility seems to be too high to be justified by fundamental shocks. Relative to a Normal distribution large exchange rate movements occur more frequently. In the extreme, single (daily) exchange rate changes can exceed the 10 percent level. Such strong reactions are not necessarily caused by fundamental shocks. Indeed, the relationship between news and exchange rate fluctuations seems to be rather fuzzy so that even lasting bubbles may evolve over time.

In section 2.2, we investigate several models which try to mimic these features. These contributions have in common that at least some of the agents are assumed to follow a rule governed behavior and that the interaction between the agents may cause complicated feedback dynamics. One important class of rules - which is neglected in standard exchange rate models - is discussed below.

2.1.3 Excursus: On Technical Analysis

Technical analysis has been a part of financial practice for many decades. Charles Dow already formulated the basic tenets of technical analysis in the late 1880s (Murphy 1999). Despite the widespread use of technical analysis, economists traditionally are very skeptical of its value. In contrast to fundamental analysis, this discipline has not received much academic scrutiny. For Lo et al. (2000), technical analysis has been an orphan from the very start. In some circles, technical analysis is known as "voodoo finance": the difference between fundamental and technical analysis is said to be not unlike the difference between astronomy and astrology. This bad reputation is partly caused by its own supporters. For instance, Feeny (1989) claims that technical analysis is more an art than a science.

The general goal of technical analysis is to identify regularities in the time series of prices by extracting (nonlinear) patterns from noisy data. Implicit in this goal is the recognition that some price movements contribute significantly to the formation of a specific pattern, while others are merely random fluctuations to be ignored. In many cases, this signal extraction is done by visual inspections of price charts. Such subjective elements naturally add to the state of controversy and confusion about technical analysis.

The plan of this excursus is as follows. Section 2.1.3.1 describes the "philosophy" of technical analysis and presents some of its most important trading rules. Note that the aim of this exercise is not to discuss technical analysis in a critical perspective, but to show how it is seen by its own supporters. Section 2.1.3.2 summarizes some survey studies about the use of technical instruments. Section 2.1.3.3 deals with the profitability of technical trading rules and section 2.1.3.4 draws some conclusions.

2.1.3.1 Technical Trading Rules

Standard manuals of technical analysis are Murphy (1999), Pring (1991), and Edwards and Magee (1966). Shorter introductions include Neely (1997), Feeny (1989), and Dunis (1989). We concentrate mainly on the contribution of Murphy. His book is regarded by many technical traders as the "Bible" of technical analysis.

2.1.3.1.1 The Philosophy of Technical Analysis

Murphy defines technical analysis as the study of market action for the purpose of forecasting future price trends. Technical analysis is based on three premises. First, market action discounts everything. Technicians believe that anything that can possibly affect the price – fundamentally, politically, psychologically, or otherwise – is actually reflected in the price of that market. A study of price action is all that is required. Second, prices move in trends. Predictable trends are essential to the success of technical analysis because they enable traders to profit by buying (selling) when the price is rising (falling). The existence of trends is often explained by appealing to Newton's law of motion. An object in motion tends to continue in motion until some external force causes it to change direction. Third, history repeats itself. Price patterns indicate the bullish or bearish psychology of the market. Since the human psychology tends not to change, the future is just a repetition of the past. Put differently, traders tend to react the same way when confronted by the same conditions.

Technical analysis is a very flexible instrument. Its supporters claim that it can handle different time dimensions. Whether the user is trading the intra-day minute by minute changes for day trading purposes or trend trading the intermediate trend, the same principles apply. For Murphy, even long range forecasting over a couple of months has proven to be profitable.

Technical trading rules can be divided into two classes: graphical chart analysis and quantitative analysis. Charting involves graphing the history of prices over some period to predict future patterns in the data from the existence of past patterns. Its advocates admit that this subjective system requires the analyst to use judgement and skill in finding and interpreting patterns. Quantitative rules impose consistency and discipline on the technician by requiring him to use rules based on mathematical functions of past prices. Next, we present some leading examples for both areas.

2.1.3.1.2 Some Graphical Tools

Technical analysis aims at participating in a trend ("the trend is your friend"). But what is meant by a trend? There are three kinds of trends. An uptrend is a situation in which each successive rally closes higher than the previous rally high, and each successive rally low closes higher than the previous rally low. In other words, an uptrend has a pattern of rising peaks and troughs. The opposite situation, with successively lower peaks and troughs defines a downtrend. If peaks and troughs evolve horizontally, the pattern is called a sideways trend.

A trough is also called support. The term indicates that support is a level where buying interest is sufficiently strong to overcome selling pressure. As a result, a decline is halted and prices turn back up again. Resistance is the opposite of support and represents a price level where selling pressure overcomes buying pressure and a price advance is turned back. Figure 9 contains an example of an uptrend.[3] Points 1 and 3 are resistance levels, and points 2 and 4 are support levels. In an uptrend, the support and resistance levels show an ascending pattern. Murphy notes that there is a tendency for round and magic numbers to stop advances or declines. These numbers act as psychological support or resistance levels.

[3] Price histories can be plotted in various ways. In the line chart, the closing prices are simply connected by a line. A bar chart represents each day's high and low prices by a vertical bar. In addition, open and closing prices are marked by ticks to the bar. Japanese candlestick charts are similar to bar charts. However, the range between open and close prices is drawn as a body. If the close is higher (lower) than the open, the body is white (black). Finally, point and figure charts visualize the price action by alternating columns of X's and O's. An X column (O) stands for rising (declining) prices. X's or O's are vertically added until the trend reverses (there is no time scale).

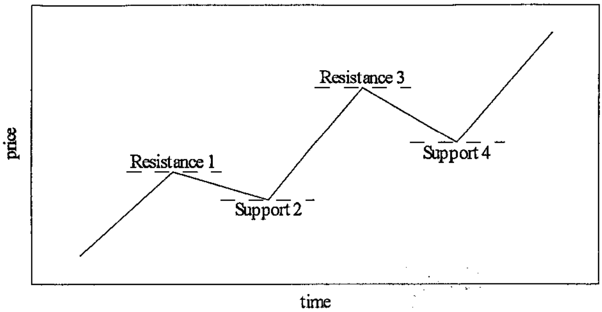

Figure 9: Support and Resistance. The solid line indicates the price evolution. Adopted from Murphy (1999).

One major tool of technical analysis are trend lines. An up trendline is a straight line drawn upward along successive reaction lows, whereas a down trendline is drawn along successive rally peaks. Sometimes prices even trend between two parallel lines - the basic trendline and the channel line. Figure 10 displays an example of a so called trend channel. The dashed line is a basic up trendline along the lows. The dotted line (channel line) connects the price peaks. Such a trend can easily be exploited. The basic up trendline can be used for building up new long positions whereas the channel line can be used for short term profit taking. More aggressive traders might even use the channel line to initiate a countertrend short position. But according to Murphy, trading in the opposite direction of a prevailing trend is dangerous.

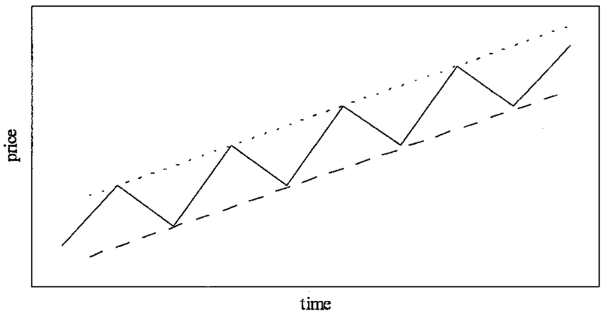

Figure 10: Trend Channel. The solid line indicates the price evolution, the dashed line the up trendline and the dotted line the channel line. Adopted from Murphy (1999).

Most changes in trend are not very abrupt affairs. In fact, important changes in trend require a period of transition, where the market moves sideways. Then, the question arises, if there will be a trend reversal or a trend continuation? Technical analysis claims that there exist reliable price patterns which indicate the future direction of the market. Next, we discuss examples of one reversal and one continuation pattern.

One of the most important reversal patterns is the head and shoulders formation. Suppose the situation of an uptrend, where a series of ascending peaks and troughs gradually begins to lose momentum (compare figure 11). One speaks of a head and shoulders pattern if the following properties are fulfilled. There are two local peaks – the shoulders – about the same height. The peak in between – the head – is higher than the shoulders. Two local minima define a neckline. Once this neckline is broken, the pattern is completed and a reversal is predicted. Clearly, the crossing of the neckline indicates a sell signal. For a downtrend, the inverse pattern holds.

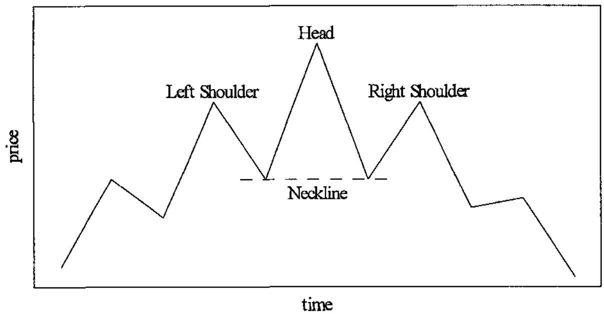

Figure 11: Head and Shoulders. The solid line indicates the price evolution. Adopted from Murphy (1999).

Continuation patterns indicate that the sideways price movement is nothing more than a pause in the prevailing trend and that the next move will be in the same direction as the trend that preceded the formation. Triangles are popular continuation patterns. They are characterized by converging trendlines. Naturally, the minimum requirement for a triangle is four reversal points. Figure 12 shows a bullish triangle with six reversal points. The three peaks and three troughs form decreasing waves. A trend is expected to resume if prices break out in its prior direction.

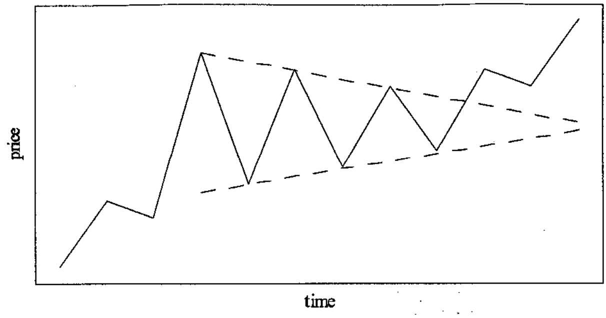

Figure 12: Triangle. The solid line indicates the price evolution, the dashed lines the triangle. Adopted from Murphy (1999).

2.1.3.1.3 Some Quantitative Tools

Chart analysis is largely subjective and therefore does not lend itself that well to computerization.[4] Quantitative rules, by contrast, can easily be programmed into a computer, which then automatically generates buy and sell signals. While two technicians may disagree as to whether a given price pattern is a triangle, quantitative signals are precise and not open to debate.

Moving averages are one of the most versatile and widely used technical indicators. For instance, a simple 10 day moving average is defined as

$$MA(t,10) = \frac{1}{10}\sum_{s=1}^{10} S_{t-s}, \tag{20}$$

where S_t is the price in period t. This rule simply averages the prices of the last 10 days. Instead of giving equal weights to each day's prices, some analysts prefer to put more weight on the most recent prices.

A moving average is essentially a trend following device. Its purpose is to signal that a new trend has begun or that an old trend has ended. More specifically, trading signals are derived as follows. When the price moves above the moving average, a buy signal is generated. A sell signal is given when the price moves below the moving average. Note

[4] However, Lo et al. (2000) have recently tried to formalize different chart patterns so that they can be detected by computer implemented algorithms.

that the shorter the moving average, the more signals are produced. If the average is too sensitive, some of the short term random price movements activate false signals. However, shorter averages have the advantage of giving trend signals earlier in the move.

Moving averages work well when the market is in a trending phase. In such a situation, Murphy recommends switching the program to automatic: let profits run! But the market is not always trending. If prices fluctuate in a horizontal price band, the class of oscillator rules should be applied.

Momentum rules are the most basic application of oscillator analysis and measure the velocity of price changes. For example, the 10 day momentum rule is defined as

$$MR(t,10) = S_t - S_{t-10}. \tag{21}$$

The so called momentum line is obtained by subtracting the price 10 days ago from the last price. For non-trending markets, the momentum line fluctuates around the zero line. Similar as in the moving average case, the time lag can be varied.

The idea behind this rule is very simple. When the oscillator reaches an extreme value, the current price move has gone too far and is due for a correction. For low values of the momentum rule, the market is said to be overbought. High values indicate that the market is oversold. Hence, buying is suggested at lower extremes and selling near upper extremes.

For some technicians, the class of mechanical trading rules has the disadvantage that it sacrifices some information that a skilled chartist might discern from the data. The moving average, for example, is said to be a follower and not a leader. It follows a market and informs that a trend has begun, but only after the fact.

To sum up, graphical tools are very flexible, but also lend themselves to subjective effects. Subjectivity can permit emotions like fear or greed to affect the trading strategy. Although quantitative tools provide clear signals, some relevant aspects may be neglected. It should be noted that both areas have a lot in common. Most importantly, they suggest trading in the direction of the trend.

2.1.3.2 The Popularity of Technical Analysis

Taylor and Allen (1992) were the first to conduct a survey study to determine the use of technical analysis in the foreign exchange market. They sent questionnaires on behalf of the Bank of England to chief foreign exchange dealers based in London. Their main findings are as follows. Chartism appears to be used most for forecasting over short time horizons (compare table 4). At the shortest horizons (intraday to 1 week), approximately 90 percent of the respondents use some chartists rules when forming exchange rate forecasts, with 60 percent judging charts to be at least as important as fundamentals. At longer horizons (6 months), the use of fundamentals increases. At the longest horizons (1 year), the skew towards fundamentals is most pronounced, with around 20 percent of respondents relying on pure fundamentals and 83 percent judging fundamentals to be more important than charts.

In addition, only 8 percent of the respondents thought the two approaches to be competing to the point of being mutually exclusive, the rest held the approaches to be complementary to a greater or lesser extent. Hence, it is not uncommon for professional traders to have two opinions at the same time regarding the future price evolution. One opinion may relate to the long-term trend which is based on fundamentals. The other opinion may be concerned with the short-term trend, which is based on the technical conditions of the market.

	0	1	2	3	4	5	6	7	8	9	10
intraday	5.6	3.4	15.6	13.4	6.7	15.6	8.4	10.6	7.3	2.8	10.6
1 week	1.8	2.5	4.9	13.5	18.4	20.2	11.0	11.0	4.3	3.1	9.2
6 months	1.8	1.8	0.9	1.8	5.5	12.8	14.7	18.3	16.5	7.3	18.3
1 year	2.4	1.2	1.2	1.2	1.2	9.6	4.8	18.1	18.1	21.7	20.5

Table 4: Technical versus Fundamental Analysis. The table contains, for each horizon, the percentage response to the question: please indicate on a scale (0=pure chartism, 10=pure fundamentalism) the relative importance you attach to technical analysis versus fundamental analysis of currencies. Source: Taylor and Allen (1992).

The results of Taylor and Allen have been confirmed for many foreign exchange markets. Compare, for instance, Menkhoff (1997) for the German market, Lui and Mole (1998) for the Hong Kong market, Cheung and Wong (2000) for the Hong Kong, Tokyo and Singapore markets, and Cheung and Chinn (2001) for the US market. Menkhoff points out that the use of technical analysis is independent from the age (which may be seen as a proxy for experience), the position or the education level of the traders. Moreover, there is no recognizable relationship between company size and the preferred class of analysis.

2.1.3.3 The Profitability of Technical Analysis

Virtually all the research on technical trading rules in the foreign exchange market concludes that these tools are successful at predicting exchange rate movements.[5] These studies can be divided into three groups. The first group (LeBaron 1999, Neely 1998, Menkhoff and Schlumberger 1995, Levich and Thomas 1993, Brock et al. 1992, Schulmeister 1988, Sweeny 1986, and Dooley and Shafer 1983) is concerned with the profitability of mechanical trading rules. A variety of the most important quantitative rules appears to be useful for different markets and time periods.

More recent work shows that graphical chart rules also display predictive power for exchange rates. This line of research faces the problem of correctly identifying patterns like the head and shoulders formation. However, the contributions of Osler (1998, 2000, 2001) and Lo et al. (2000) support the view that chartism can generate profits.

All these studies are based on a range of rules chosen ex post. For this reason, there remains some doubt as to whether the reported profits could have been earned by a trader who had to make a choice about what rules to use at the beginning of the sample period. The third area of research thus employs search procedures such as genetic algorithms, genetic programming and artificial neural networks to identify optimal trading rules. Examples of this group include Skouras (2001), Fernandez-Rodriguez et al. (2000), and Neely et al. (1997) who find strong evidence for excess returns.

[5] One of the few exceptions is Curcio et al. (1997).

2.1.3.4 Summary

The term technical analysis is a general heading for a myriad of trading strategies. There are probably as many methods of combining and interpreting the various techniques as there are chartists themselves. However, all technical traders subscribe to the notion of market momentum and rely on some sort of feedback rule. They employ either graphical or mechanical rules to clarify the direction of market movements. Although different, these practices undoubtedly share some similar (extrapolative) elements.

Empirical work surprisingly reveals that at least some technical trading rules have been profitable in the past. Hence, whatever the reasons for this may be, agents who incorporate technical analysis into their investment decision process should not be called irrational per se. But even if one doubts this finding, one cannot ignore that the large majority of foreign exchange professionals relies on both technical and fundamental analysis. At shorter horizons, there exists a skew towards reliance on technical analysis as opposed to fundamental analysis, but the skew becomes steadily reversed as the length of the forecast horizon is extended. Since the trading volume of the foreign exchange market mainly reflects very short-term speculative transactions, one cannot rule out that technical trading itself has an impact on the price formation process. This aspect is crucial for the models we survey in section 2.2.

2.2 Survey of the Literature

This section surveys three related fields of research in behavioral finance: the noise trader approach, chartists-fundamentalists models and agent-based computation. Although substantially different in methods and style, these emerging areas are all attempts to go beyond economic theories based on fully rational agents and market efficiency. In particular, the noise trader approach studies the impact of non-fully rational traders on the price formation when arbitrage is limited. Chartists-fundamentalists models are more specific and focus on the interaction of a limited set of popular trading rules. Agent-based computation models are designed to capture complex learning behavior and dynamics. In general, these models do apply to all financial markets.

2.2.1 The Noise Trader Approach

The noise trader approach is a promising attempt to offer an alternative to the efficient market hypothesis (Menkhoff 1998, Frankel and Rose 1995). The noise trader approach, summarized by Shleifer and Summers (1990), rests on two core assumptions. First, some investors are not fully rational. Their demand for risky assets is affected by beliefs or sentiments that are not fully consistent with economic fundamentals. Second, arbitrage – defined as trading by fully rational investors not subject to such sentiments – is risky and therefore limited. Together, these assumptions imply that shifts in investor sentiment are not fully countered by arbitrageurs and so affect asset prices.[6]

Obviously, market participants are supposed to be heterogeneous. There are at least two types of investors. Arbitrageurs (also called rational speculators) form fully rational expectations about asset returns whereas the opinions and trading patterns of noise traders (or liquidity traders) are subject to systematic biases.

The noise trader approach aims at explaining some well known anomalies of financial markets and may also help understanding some of its broad features like the widespread use of simple heuristic investment strategies. Moreover, the noise trader approach may allow different normative conclusions than its main competitor, the efficient market approach. Next, we discuss the noise trader approach in more detail.

2.2.1.1 The Limits of Arbitrage Assumption

Arbitrage plays a central role in standard finance theory. As long as an asset has a perfect substitute, the price of the asset equals the price of the substitute portfolio. For example, if the price of an asset falls below that of a substitute portfolio, arbitrageurs simultaneously sell the portfolio and buy the asset until the prices are equalized. Such textbook arbitrage requires no capital and entails no risk. Since arbitrageurs have perfectly elastic demand for the asset, arbitrage brings prices to fundamental values and keeps markets efficient. However, not every asset has a perfect substitute.

[6] The main features of the noise trader approach are also described by Black (1986). According to Piron (1991), the basic ideas go back to Keynes (1936).

Then, two types of risk limit arbitrage (Shleifer and Summers 1990): fundamental risk and noise trader risk. Fundamental risk represents the uncertainty connected with the evolution of the fundamental price. For instance, selling overvalued stocks is risky because there is always a chance that the market will do very well (good news). The second type of risk comes from the unpredictability of the future resale price introduced by the noise traders. If future mispricing is more extreme than when the arbitrage trade is put on, the arbitrageur suffers a temporary loss on his position. The fear of such a loss limits the arbitrage positions. Prices are not driven to fundamentals.

The latter point is made more explicit in Shleifer and Vishny (1997). Note first that even simple arbitrage operations are more complex than the textbook definition suggests. Arbitrage always requires capital (e.g. the initial margin for a futures contract) and is risky (e.g. due to differences in trading hours). Commonly, arbitrage is conducted by relatively few professional, highly specialized investors who combine their knowledge with resources of outside investors to take large positions. In the words of Shleifer and Vishny, brains and resources are separated by an agency relationship. The money comes from wealthy individuals or institutions such as banks or pension funds with only limited knowledge of individual markets, and is invested by arbitrageurs with highly specialized knowledge of these markets. If arbitrageurs could act independently, they would generally be more aggressive when prices move further from fundamental values. But when the capital owner do not know what the manager is doing, they will only observe him losing money. They may come to the conclusion that the manager is not as competent as they previously thought and may withdraw their money. In this sense, arbitrageurs can become most constrained when they have the best opportunities. The fear of this scenario makes them more cautious when they put on their initial trades. Risk averse arbitrageurs may even voluntarily liquidate their positions to prevent the need to liquidate the portfolio under pressure. Resources dedicated to long-term arbitrage against fundamental mispricing are scare.

Shleifer and Summers (1990) give a prominent real-life example to illustrate the risk associated with arbitrage. Japanese equities in the 1980s have been traded at price earning ratios between 20 and 60. Expected growth rates of dividends and risk premia required to justify such ratios seem unrealistically high. Nonetheless, an investor who believes that Japanese equities are overvalued and wants to sell them short faces two types of risk. First, what if the Japanese economy actually does perform so well that these prices are justified (the fundamental risk). Second, how much more out of line can prices get, and for how long, before Japanese equities return to more realistic prices (the noise trader risk)? Investors who sold Japanese stocks short in 1985, when the price earnings ratio was 30, made heavy losses as the ratio rose to 60 in 1986.

2.2.1.2 The Investor Sentiment Assumption

Of course, some shifts in investor demand for assets are fully rational. For instance, investors naturally adjust their positions due to public news announcements. But the demand also reacts in response to changes in expectations or sentiments that are not completely justified by information. Shleifer and Summers (1990) identify two kinds of sources for shifts in investor sentiment: the belief in pseudo signals and the use of popular models.

An example for a pseudo signal, also called noise, may be the advice of a financial guru. Black (1986) uses the term noise in the sense of a large number of small events. For Black, noise is often a causal factor much more powerful than a small number of large events can be. He points out that noise trading is essential to the existence of liquid markets. If there is no noise trading, there will only be very little trading. The argument is that a person who wants to trade needs another person with opposite beliefs. To explain the high trading volume in the financial markets, it is not reasonable to assume that differences in beliefs are merely the result of different information. Investors act on noise as if it were information.

After Shiller (1990), the expression popular models refers to the models that are used by the broad masses of economic actors to form their decisions. Naturally, these differ to those held by the economists. Shiller presents survey data on the use of popular models during the stock market crash of 1987. Both (wealthy) individual and institutional investors were asked about the causes of the crash. The respondents had the choice between fundamental reasons and investor psychology. Surprisingly, 2/3 of the US investors and 3/4 of the Japanese investors selected the latter one. Moreover, 1/3 of the investors said that they anticipated the crash. Asked about the methods with which they came up with this forecast, they wrote intuition, gut feeling, or named extremely simple models like trend reversals. Finally, the questionnaire reveals that around 40 percent of the investors experienced symptoms of anxiety like sweaty palms, tightness in the chest or rapid pulse. Shiller concludes that a crash may be caused by people reacting to each other with heightened attention and emotion, trying to fathom what other investors were likely to do, and falling back on intuitive models of price reversal or continuation. No recognizable exogenous trigger has to appear for a crash! Such popular models may create feedback systems with complicated dynamics. De Bondt (1998) presents similar evidence on the behavior of small individual investors. Many people discover naive price patterns in past price movements and share popular models of value.

Note that shifts in investor sentiment only matter if they are correlated across traders. If all investors trade randomly, their positions cancel out. But there exists a lot of empirical evidence that trading activity based on pseudo signals and popular models is correlated and thus leads to aggregated demand shifts. As is demonstrated in section 2.1.3, chart analysis has a large subjective element: there are probably as many methods of combining and interpreting the various techniques as there are chartists themselves. But most chartists subscribe to the notion of market momentum and rely on some sort of feedback rule (Murphy 1999). These practices undoubtedly share some similar elements.

In addition, psychological experiments show that judgment biases afflicting investors in processing information tend to be the same. Two important phenomena documented by psychologists are the conservatism and representativeness heuristic. Conservatism states that individuals are slow to change their beliefs in the face of new information. The second heuristic describes the tendency of individuals to view events as typical or representative of some specific class and to ignore the laws of probability. For example, people often think they see patterns in truly random sequences. Barberis et al. (1998) develop a model of investor sentiment explicitly based on these regularities which explains both underreaction and overreaction to news.

People who regularly interact with each other tend to think and behave similarly. According to Shiller (1995), there are two main strands of literature dealing with herding behavior. In the models of information cascades (Banerjee 1992), people acquire information in sequences by observing the actions of other individuals who precede them in the sequence. Consider, for example, a situation where a group of agents has to choose between two restaurants that are next to each other. Each agent has an imperfect signal about the quality of the restaurant. Suppose further that the first individual follows his signal. Those who come after may rationally ignore their own signal, deciding that these signals are dominated by the information revealed by their predecessors′ decisions to go to one of the restaurants. In the end, everyone may end up in the wrong restaurant, if the first signal is bad. Transferred to financial markets, the trading activity of one investor conveys information to other investors that can cause the latter to react, leading to a cascade of trading. A small piece of news can result in a large change in price.

The second field in concerned with the interpersonal communication of the market participants. Although informational cascades appear to be important, it seems that conversation is more relevant to describe mass behavior, especially in financial markets. Conversation not only serves to exchange information, it also reinforces memories of pieces of information. If a certain "theory" becomes established, it may persist for a long time (think of fads or fashions). Such failures in human judgment generate parallel behavior (Shiller 1995).

2.2.1.3 The Importance of Noise Traders

From these comments, there should remain little doubt that noise traders influence the demand for assets. But can this group be big enough to have a permanent impact on the price formation? According to Friedman (1953), these effects can be neglected. Rational speculation is always stabilizing. Investors who are engaged in destabilizing speculation – selling when the currency is low in price and buying when it is high – lose their money and drop out the market. Therefore, noise traders might learn from their errors and convert to arbitrageurs as time proceeds.

However, De Long et al. (1990a) demonstrate that if a market is dominated by investors who follow positive feedback strategies (buy when prices rise and sell when prices fall) it need no longer be optimal for arbitrageurs to counter shifts in the demand of these investors. Instead, it may pay for arbitrageurs to jump on the bandwagon themselves. Arbitrageurs then optimally buy the assets that positive feedback traders get interested in when their prices rise. When prices increase, fed by the buying of other investors, arbitrageurs sell out near the top and take their profits. The effect of arbitrage is to stimulate the interest of other investors, thus contributing to the movement of prices away from fundamentals. Although arbitrageurs eventually sell out and help prices return to fundamentals, in the short run they feed the bubble rather than help it to dissolve. Rational speculation is not stabilizing per se (similar Baumol 1957).

The point that noise traders lose money and eventually disappear is also not self-evident. De Long et al. (1990b) show that noise traders may earn a higher expected return than rational arbitrageurs. Their argument is that noise traders enlarge price volatility. Facing additional risk, rational speculators with sufficient risk aversion will limit their positions against noise traders. If noise traders underestimate the risk, they invest more in the risky asset than risk-averse arbitrageurs. Clearly, noise traders earn higher average returns for bearing self created risk and not for bearing more fundamental risk. They profit from their own destabilizing influence. This implies that noise traders as a group don't necessarily have to disappear from the market. However, there might be a high dispersion. Some noise traders end up very rich whereas the others lose their capital.

Thus, learning and imitation may not adversely affect noise traders. When noise traders earn high average returns, many other investors might imitate them, ignoring the fact that they took more risk and just got lucky. Such imitation increases the impact of noise traders. As noted by Shleifer and Summers (1990), noise traders who do well become more aggressive since they tend to attribute their investment success to skill rather than luck. Finally, one should not forget that new investors enter the market all the time and old investors who have lost money come back. These investors are subject to the same judgment biases as the current survivors in the market, thereby adding to the effect of judgement biases on demand.

2.2.1.4 Results and Implications

When arbitrage is limited and investor demand for assets responds to noise and to predictions of popular models, prices move in response to these changes in demand as well as to changes in fundamentals. Arbitrageurs counter the shifts in demand prompted by changes in investor sentiment, but do not completely eliminate the effects of such shifts on the price.

From a dynamical perspective, Black (1986) sketches a possible price formation process as follows. With a lot of noise traders in the market, it pays for rational investors to seek out costly information which they will trade on. The farther the price of an asset diverges from its value, the more aggressive the rational traders will become. More of them will come in, and they will take larger positions. Eventually, the price will tend to move back toward its value over time. If the adjustment occurs quickly, technical traders will perceive it and speed it up. In general, the rational traders will not take large enough positions to eliminate the noise completely. For one thing, taking larger positions means taking more risk, and the rational traders can never be sure that they are trading on information rather than on noise.

The noise trader approach is obviously able to explain some stylized facts of financial markets. Prices are more volatile than is warranted by changes in fundamentals since they respond to shifts in investor sentiment. If price movements triggered by news are prolonged by positive feedback trading, asset prices overreact to news. The mispricing of an asset may even culminate in a large bubble with a consecutive crash. The interaction between noise traders and arbitrageurs may cause patterns like mean reversion. For instance, positive feedback strategies may extend price trends and hence lead to a positive autocorrelation of returns at short horizons. Eventual return of prices to fundamentals entails a negative autocorrelation at longer horizons.

The noise trader approach raises two important normative questions. First, should something be done to prevent noise traders from suffering from their errors? Some noise traders may earn higher average returns because they bear more risk. But for a lot of traders speculation is hurtful. Shleifer and Summers (1990) argue that the case of making it costly for investors to bet on financial markets to protect them is comparable to the case of prohibiting casinos or state lotteries. We do not share this point of view. Roulette is obviously a gamble, its rules and chances are easy to understand. Financial markets are much more complicated. Market participants should at least be informed about the risk connected with these speculative markets. Interpreting stock markets as sound long-term investment opportunities may result in social disasters.

Second, do noise traders impose a cost on the rest of the market participants? The overall impact of noise trading on the society can be negative. For instance, noise trading makes returns on assets more risky, and thus can reduce physical investment. In the foreign exchange market, noise trading can distort the flow of goods between countries and lead to an inefficient choice of production. In addition, noise trading may force managers to focus on the short-term and to bias the choice of investments against long-term projects. We would like to add that a huge stock market crash may even trigger a recession. A sharp decline in asset prices, as for example observed in Japan in the 1990s (Bayou 2001, Westerhoff 1999), goes along with a wealth reduction and thus depresses consumer confidence. Lower expenditures naturally slow down growth. A stock market crash may also draw the whole financial system into a crisis. Remember that more than 10 years after the burst of the asset price bubble, the Japanese banking sector still is in deep trouble. However, for Shleifer and Summers (1990) the issue of what should be done to reduce the costs is open. The consequences of measures like short-term capital gains taxes or transaction costs still need a more systematic evaluation.

2.2.1.5 Summary

It is no question that forces of noise trading are present in financial markets. Their existence is all too obvious. But to which extent do these forces hold in a broader empirical investigation? A first answer to this question is given by a survey study of Menkhoff (1998). Menkhoff sent questionnaires to foreign exchange professionals from banks and fund management companies trading in Germany in 1992. The questions were especially designed to test the core elements of the noise trader approach. In general, the evidence gained is in strong support with the assumptions of the noise trader approach.

His main results are: arbitrage is limited. The endurance to hold on to loss positions, which might turn profitable in the long run, is usually limited to some months. Often, investors have to close their loss positions within weeks. Agents believe in the importance of psychology for asset pricing. Their demand is influenced by pseudo signals and popular models. Arbitrageurs and noise traders are not as different as expected. Almost all respondents rely at least partially on non-fundamental popular models. Menkhoff explains this apparent contradiction in the noise trader approach with the finding of De Long et al. (1990a). Rational arbitrageurs might try to exploit noise traders while jumping on the bandwagon in an early stage. Hence, investors rationally may use feedback rules to amplify a price trend so that noise traders automatically start to trade in the same direction. For Menkhoff, this understanding does not sacrifice the distinction between rational and not fully rational traders.

We do not agree with this argument. Our understanding of the empirical evidence is that a clear distinction between rational and irrational traders is not appropriate. Our perspective is that all agents are boundedly rational. This term goes back to Simon (1955) and recognizes the cognitive limitations of a decision maker with respect to both knowledge (information and theory!) and computational capacity. The bounded rationality approach is strongly supported by psychological experiments (compare the survey of Conlisk 1996, Camerer 1995, Slovic, Lichtenstein and Fischhoff 1988, and Kahneman, Slovic and Tversky 1986). They indicate that individuals tend to simplify decision problems and use approximating procedures rather than investigating all possible alternatives to find the optimum. The latter would be too complicated for the agents or even impossible. The chartists-fundamentalists literature follows this aspect.

2.2.2 Chartists and Fundamentalists

The class of chartists-fundamentalists models is based on the observation that agents strongly rely on simple technical and fundamental trading rules when determining their speculative investment positions (Lui and Mole 1998, Menkhoff 1997, Taylor and Allen 1992). Chartists-fundamentalists models may be understood as a special branch of the noise trader approach. Since these models aim at replicating realistic asset price fluctuations, the style of analysis is less analytically but more numerically oriented. Next, a selection of the most influential chartists-fundamentalists models is discussed. The aim is to distinguish between various mechanisms and setups which are able to generate complex dynamics. Thus, we do not analyze the whole model but concentrate on its most important features.

2.2.2.1 The Model of Frenkel

Frenkel (1997), expanding on work by Cutler et al. (1990), combines different types of speculators to account for heterogeneous expectations. Frenkel applies a simple asset market approach. The exchange rate is determined by the stock of domestic assets relative to foreign assets and the demand for foreign assets relative to domestic assets

$$s_t = m_t + d_t + u_t, \qquad (22)$$

where s_t denotes the Log of the spot rate at time t, m_t the Log of the relative supply of domestic assets, d_t the Log of the relative demand for foreign assets, and u_t some random shocks. For the sake of simplicity, the supply of assets is assumed to be constant so that only the relative demand and random factors influence the exchange rate.

The model consists of two types of traders: chartists and fundamentalists. The relative demand of chartists for foreign assets is

$$d_t^C = a^1(s_t - s_{t-1}) + a^2(s_{t-1} - s_{t-2}), \qquad (23)$$

where a^1 and a^2 are positive reaction coefficients. According to (23), the relative demand of the chartists is based on a linear trend. Fundamentalists assume that the exchange rate moves towards its fundamental value f

$$d_t^F = b(f - s_t), \qquad (24)$$

where b is a positive reaction coefficient.

The total demand of both groups is given by

$$d_t = wd_t^C + (1-w)d_t^F. \qquad (25)$$

Note that the market shares of chartists w and fundamentalists (1-w) are constant.

Combining (22)-(25) yields the solution of the model

$$s_t + \frac{w(a^1 - a^2)}{1 - wa^1 + (1-w)b} s_{t-1} + \frac{wa^2}{1 - wa^1 + (1-w)b} s_{t-2} = \frac{m + (1-w)bf + u_t}{1 - wa^1 + (1-w)b}, \qquad (26)$$

which is a two-dimensional linear difference equation (u=0). It is well known that such a system can generate oscillations.

Frenkel analytically shows that fundamentalists dampen price fluctuations: the higher the share of the fundamentalists, the lower the volatility of the exchange rates. In addition, the critical value which ensures stability rises. The model also replicates some empirical regularities of the foreign exchange market. Due to positive feedback trading, the exchange rate overreacts to a random shock. In the long run, however, the exchange rate converges towards its fixed point. This implies that returns must be negatively correlated over some horizons (mean reversion).

2.2.2.2 The Model of Day and Huang

Day and Huang (1990) and Huang and Day (1993) develop a theory that explains the tendency of stock markets to exhibit alternating periods of bear and bull markets. Their model incorporates three types of market participants: chartists, fundamentalists and market makers.

In Huang and Day (1993), the demand of the fundamentalists is a piece-wise linear function. Let S^B and S^T be bottoming and topping prices, and let S' and S'' be threshold buying and selling prices, with $S^B < S' < F < S'' < S^T$. The parameters S' and S'' determine when fundamentalists are active or inactive. F is the fundamental price. The excess demand of the fundamentalists is assumed to be

$$d_t^F = \begin{cases} A, & S_t \le S^B \\ b(S' - S_t), & S^B \le S_t \le S' \\ 0, & S' \le S_t \le S'' \\ -b(S_t - S''), & S'' \le S_t \le S^T \\ -A, & S^T \le S_t \end{cases} . \qquad (27)$$

When S'<S<S'', fundamentalists hold their positions, when S>S'', they enter the market to sell, and when S<S′, they enter the market to buy. The positive risk coefficient b is constant.[7]

Chartists believe that the price indicates all relevant information about the future development of the market. If $S<F$, a bear market is expected, and if $S>F$, a bull market is expected. Their excess demand is given as

$$d_t^C = a(S_t - F), \qquad (28)$$

where a is a positive reaction coefficient.

Market makers mediate transactions, i.e. they set the price in response to excess demand or supply. Market makers supply excess demand from their inventory or accumulate inventory when there is an excess supply. Formally, they determine the change in price by a continuous monotonically increasing function of excess demand. Huang and Day assume that

$$S_{t+1} = S_t + c(d_t^F + d_t^C), \qquad (29)$$

where c is a positive price adjustment coefficient.

Inserting (27) and (28) in (29) yields a one-dimensional nonlinear difference equation. Relative to linear models, where the fluctuations are derived from stochastic shocks rather than from intrinsic market interaction, the models of Day and Huang (1990) and Huang and Day (1993) generate deterministic chaotic motion. The interaction between chartists and fundamentalists endogenously causes an irregular switching between bear and bull markets.[8]

2.2.2.3 The Model of de Grauwe, Dewachter and Embrechts

De Grauwe et al. (1993) focus on the foreign exchange market. Their starting point for analyzing exchange rate dynamics is

[7] Day and Huang (1990) assume a monotonically decreasing excess demand function of fundamentalists, which is falling rapidly near S^B, flattening out near F (for $S=F$, the demand is zero) and then falling rapidly again near S^T. Overall, this function yields similar results.

[8] A similar model is constructed by Chiarella (1992). In his model, the excess demand of the fundamentalists is a linear function. Chaos results from a nonlinear excess demand function of chartists.

$$S_t = X_t E_t[S_{t+1}]^a, \tag{30}$$

where X_t is a reduced form equation describing the structure of the economy relevant for the exchange rate. To simplify the analysis, X_t is assumed to be constant and normalized to 1.[9] Hence, the exchange rate in period t depends solely on the expected future exchange rate, discounted by a ($0<a<1$).

Market participants consist of chartists and fundamentalists. Chartists extrapolate recent observed exchange rate changes into the future by using a univariate time series model

$$E_t^C[S_{t+1}/S_{t-1}] = g(S_{t-1},...,S_{t-n}), \tag{31}$$

where n indicates the lag length of the applied rule.

Fundamentalists have regressive expectations

$$E_t^F[S_{t+1}/S_{t-1}] = (F_{t-1}/S_{t-1})^b. \tag{32}$$

The fundamentalists expect the exchange rate to return to its fundamental value F at the speed b during the next period ($0<b<1$).

The size of both groups remains constant over time. However, the chartists are a homogenous group whereas the fundamentalists are characterized by heterogeneous expectations. This novel idea introduces an important nonlinearity into the model. Suppose that among a group of fundamentalists, each makes a forecast of the fundamental value, and that these estimates are normally distributed around the true fundamental value. When the exchange rate is equal to the fundamental value, half the fundamentalists view the exchange rate as overvalued and half as undervalued. This difference in opinion makes fundamentalists net demand zero, so that they have no net effect on the market. But the more the exchange rate deviates from its fundamental value, the more influential the fundamentalists become.

Therefore, the expectations are aggregated as

$$E_t[S_{t+1}/S_{t-1}] = E_t^C[S_{t+1}/S_{t-1}]^{m_t} E_t^F[S_{t+1}/S_{t-1}]^{1-m_t}, \tag{33}$$

[9] In continued work, X_t is specified by detailed structural exchange rate models. For instance, de Grauwe and Dewachter (1992, 1993) employ a Dornbusch-type exchange rate model and da Silva (1999) a "new open economy macroeconomics model". In these environments, the dynamics can display chaotic motion.

where the market impact of the chartists is approximated as

$$m_t = 1/(1 + c(S_{t-1} - F_{t-1})^2). \tag{34}$$

The influence of the fundamentalists is given as $(1-m_t)$. The parameter c measures the degree of divergence of the fundamentalists′ estimate of F.

The model is solved by substituting (31), (32) and (34) into (33), and (33) into (30). The solution is a n+1-dimensional nonlinear difference equation. Simulations show that the model has the potential to produce chaotic motion. Note that the dynamics live from the fact that the fundamentalists are heterogeneous. Whenever $S \approx F$, fundamentalists cancel out each other so that the feedback trading of the chartists destabilizes the market.

2.2.2.4 The Model of Frankel and Froot

The chartists-fundamentalists literature basically starts with Frankel and Froot (1986, 1990).[10] They propose a model of speculative bubbles to understand the exchange rate path of the US dollar in the 1980s. The model features three groups of players: fundamentalists, chartists and portfolio managers. According to Frankel and Froot, portfolio managers are those actors who actually buy and sell foreign assets. They form their expectations as a weighted average of the predictions of the fundamentalists and the chartists. The portfolio managers update the weights over time, according to whether the fundamentalists or the chartists have recently been doing a better job of forecasting.

More precisely, Frankel and Froot also start with a general model of exchange rate determination

$$s_t = aE^P[\Delta s_{t+1}] + x_t, \tag{35}$$

where s_t is the Log of the spot rate, $E^P[\Delta s_{t+1}]$ is the rate of change in the exchange rate expected by the portfolio managers, a a positive discount factor, and x stands for other determinants.

[10] However, there exist some earlier contributions. Already in 1974, Zeeman (1974) uses the terms chartists and fundamentalists. Beja and Goldman (1980) explore a market maker model with linear demand functions of chartists and fundamentalists. Unfortunately, their work did not receive much attention. One explanation for this may be that the study of complex nonlinear dynamics was almost impossible until fast computers became available.

The expectations of the portfolio managers are a weighted average of the predictions of fundamentalists and chartists

$$E^P[\Delta s_{t+1}] = w_t E^F[\Delta s_{t+1}] + (1-w)E^C[\Delta s_{t+1}], \tag{36}$$

where $E^F[\Delta s_{t+1}]$ and $E^C[\Delta s_{t+1}]$ are the predictions of fundamentalists and chartists, and w_t is the weight given to the fundamentalists.

The fundamentalists expect the exchange rate to regress to its fundamental value

$$E^F[\Delta s_{t+1}] = b(f - s_t), \tag{37}$$

where f is the Log of the fundamental and b the speed of regression of s_t towards f. The forecast of the chartists is a random walk

$$E^C[\Delta s_{t+1}] = 0. \tag{38}$$

If fundamentals remain fixed, (35) together with (36), (37) and (38) can be rewritten as

$$\Delta s_t = ab\Delta w_t(f - s_{t-1})/(1 + abw_t). \tag{39}$$

Hence, the exchange rate path critically depends on the evolution of w_t. Suppose that the weight of the fundamentalists decreases. Then, the exchange rate moves away from its fundamental.

Frankel and Froot assume that the portfolio managers update the weight given to the predictions of the fundamentalists as

$$\Delta w_t = c(\omega_{t-1} - w_{t-1}), \tag{40}$$

where ω_{t-1} is the ex post computed weight that would have accurately predicted the change in the exchange rate in the previous period. This weight is defined as

$$\omega_{t-1} = \Delta s_t /(b(f - s_{t-1})). \tag{41}$$

Combining (40) and (41) yields

$$\Delta w_t = c(\Delta s_t /(b(f - s_{t-1})) - w_{t-1}). \tag{42}$$

The positive coefficient c controls the adaptiveness of w_t. If the fundamentalists predict well ($\Delta s_t \approx b(f - s_{t-1})$), their influence tends to rise. Since the portfolio manager can at most take one view or the other exclusively, w_t is, of course, bounded between one and zero.

Since the solution of the model has no closed analytic form, Frankel and Froot proceed with a numerical investigation. They find that if portfolio managers are slow learners, the spot rate tends to move away from its long-run equilibrium if it is perturbed (for small c the fixed point is unstable). The weight attached to the predictions of fundamentalists is falling over time (since their forecasts are bad), whereas chartists are gaining prominence. The exchange rate begins to trace out a bubble path until, eventually, the expectations of portfolio managers are determined only by the view of chartists. At this point, the bubble dynamics die out (Δw_t=0 implies Δs_t=0). Frankel and Froot conclude that their model is able to explain the large and sustained dollar appreciation in the 1980s.

2.2.2.5 The Model of Kirman

Kirman (1991, 1993) relates changes in market opinion to speculative bubbles. The traders either form opinions in a fundamental or in a chart-technical way. The opinions are influenced by the stochastic interaction between the traders. The traders talk to each other, try to determine what the market opinion is and then act accordingly.

Similar as before, the exchange rate depends on the average forecast of the agents and some other determinants

$$S_t = a(w_t E^F[\Delta S_{t+1}] + (1 - w_t) E^C[\Delta S_{t+1}]) + X_t, \quad (43)$$

where the fundamentalists hold regressive expectations

$$E^F[\Delta S_{t+1}] = b(F - S_t), \quad (44)$$

and the chartists use a linear trend forecast

$$E^C[\Delta S_{t+1}] = c(S_t - S_{t-1}). \quad (45)$$

The discount factor a and the reaction coefficients b and c are positive.

Kirman models the interaction between the traders as follows. There are two prevalent opinions in the world. Let opinion 1 be a fundamentalist view and opinion 2 a chartist view. Each of the N traders holds exactly one of these views. The state of this system, defined by the number k of agents holding view 1, evolves in two steps. First, two individuals meet at random. The first is converted to the second's opinion with probability δ. Kirman gives two arguments for such a recruitment. On the one hand, an agent may persuade another of the superiority of his choice either because of better information or a better knowledge of the functioning of the market. On the other hand, one may interpret this phenomenon as herding behavior. Trader one just imitates the behavior of trader two. In addition, with probability ε the first trader changes his opinion independently before meeting another one. This self-conversion may be thought of as the replacement of an existing trader by a new one. A probability $\varepsilon>0$ is needed to prevent the process stuck at $k=0$ or $k=N$.

Second, the agents try to assess what the majority opinion is. At time t the proportion of agents who hold a fundamentalists view is given by $q_t=k_t/N$. Each agent i observes q_t but with some noise, i.e. $q_{i,t}=q_t+\mu_{i,t}$, where $\mu\sim N(0, \sigma)$. If agent i observes that $q_{i,t}\geq 0.5$, then he acts as a fundamentalist, otherwise as a chartist. The proportion of agents who behave as fundamentalists is thus given by

$$w_t = \#\{i \mid q_{i,t} \geq 0.5\}/N, \quad (46)$$

where # stands for "the number of".

Depending on the values of the underlying parameters, the system switches between its states (k=1, 2,..., N). This process is intrinsically dynamic. There is no convergence to any particular state, but every state is always revisited. Kirman shows that the smaller ε to δ, the more time the system spends at the extremes. Clearly, if a certain number of agents act as chartists, the exchange rate is driven away from its fundamental. However, since opinions switch from time to time, fundamentalists eventually control the market and bring the exchange rate back to its long-run equilibrium.

2.2.2.6 The Model of Brock and Hommes

In the framework of Brock and Hommes (1997a, 1997b, 1998, 1999), agents adopt their beliefs over time by choosing from a finite set of different predictors. Each predictor is a function of past observations. The agents are boundedly rational in the sense that they tend to rely on predictors which have generated the highest performance in the recent past. The performance of the predictors is generally known.

Brock and Hommes (1997b, 1998, 1999) consider a standard asset pricing model with one risk-free asset and one risky asset. The risk-free asset pays a fixed rate of return R. The risky asset is traded at the price S_t and yields an uncertain dividend Y_t. The wealth in period t+1 is

$$W_{t+1} = RW_t + (S_{t+1} + Y_{t+1} - RS_t)z_t, \tag{47}$$

where z_t denotes the number of shares purchased at time t. Each investor type h is a myopic mean-variance maximizer

$$\max_{z_{h,t}} E_{h,t}[W_{t+1}] - (a/2)V[W_{t+1}] \ , \tag{48}$$

where a is the risk aversion. For simplicity, the agents share the same belief about the constant variance, i.e. $V[S_{t+1}+Y_{t+1}]=\sigma^2$. Hence, the demand of shares of trader type h is

$$z_{h,t} = E_{h,t}[S_{t+1} + Y_{t+1} - RS_t)/a\sigma^2. \tag{49}$$

In the case of zero supply of outside risky assets, market equilibrium implies

$$S_t = (\sum_h n_{h,t} E_{h,t}[S_{t+1} + Y_{t+1}]) / R, \tag{50}$$

where $n_{h,t}$ denotes the fraction of traders using predictor h at time t. Agents are assumed to prefer predictors with a high past performance. The fitness of a predictor U is generally evaluated as a weighted sum of risk-adjusted profits π

$$U_{h,t} = \pi_{h,t} + \eta U_{h,t-1}, \tag{51}$$

where η indicates the memory of the traders. The larger η, the more the fitness of a predictor depends on past profits.

Brock and Hommes discuss various simple belief types and performance measures. Overall, complex dynamics can emerge. Suppose that agents can either buy, at small but positive information costs, a sophisticated predictor H1 (say a forecast based on fundamentals) or freely obtain another simple predictor H2 (say a trend extrapolation). Predictors which yield high net profits are preferred. If all agents use the simple predictor, prices diverge from the long-run equilibrium. Since this increases the uncertainty in the market, the number of traders who are willing to pay some information costs to get the prediction H1 rises. At some time, prices are pushed back towards the long-run equilibrium and remain there for a while. With prices close to the long-run equilibrium, the prediction error of H2 becomes smaller. Hence, agents switch their beliefs to predictor H2 again, and the story repeats. Brock and Hommes conclude that there is one "centripetal force" when most agents use the sophisticated predictor and another "centrifugal force" when most agents use the simple predictor. Under some conditions, the interaction between the two forces produces chaotic motion. Adding some noise to the dynamics (Gaunersdorfer and Hommes 2000), the model is able to mimic the behavior of financial data, creating, for instance, volatility clustering.

2.2.2.7 The Model of Lux

Lux (1995, 1997, 1998) studies the economic and social interactions of speculators. On the one hand, traders switch between chartists and fundamentalists behavior because of performance differentials. On the other hand, the social interactions between the traders lead to a herding behavior. Resulting imbalances between demand and supply are managed by market makers. In a broader sense, Lux combines the economic reasoning of Frankel and Froot, and Brock and Hommes with the social reasoning of Kirman.

His model is as follows. The market makers absorb the excess demand of the chartists and fundamentalists and quote the new price

$$S_t = S_{t-1} + a(d_t^C + d_t^F), \qquad (52)$$

where a denotes the price adjustment speed.

The fundamentalists attempt to realize profits in the presence of overvaluation or under-valuation of an asset. Such a strategy implies buying (selling) when prices are below (above) the fundamental value. The excess demand of the fundamentalists is

$$d_t^F = n_t^F b(F - S_t), \qquad (53)$$

where n^F is the number of fundamentalists and b a positive reaction coefficient.

The group of the chartists is heterogeneous in the sense that chartists are either bullish or bearish. Chartists who perceive the market as bullish buy g^C units of the asset, and those who fear a negative price trend sell g^C units of the assets. The net excess demand of the chartists is

$$d_t^C = (n_t^{C+} - n_t^{C-})g^C, \tag{54}$$

where n^{C+} and n^{C-} are the numbers of traders who are optimistic and pessimistic, respectively.

In total, there are $N=n^{C+}+n^{C-}+n^F$ traders. The dynamics of the model depend on the agents′ changes of behavior and on the price movements resulting from their trading. Central to the work of Lux are two core elements. First, chartists switch between an optimistic and a pessimistic subgroup. These re-evaluations of expectations occur under the influence of the majority opinion, as well as the observed price trend. The majority opinion is formally captured by constructing an opinion index $X=(n^{C+}-n^{C-})/N$, ranging from -1 (uniform pessimistic) to +1 (uniform optimistic). Both a prevailing optimistic sentiment among the chartists and a positive price trend ($S_t-S_{t-1}>0$) favor the transition from depressed to optimistic chartists.

The second element is the switching of agents between the chartist and fundamentalist group. The traders meet individuals from the other group, compare (myopic) excess profits from both strategies and change to the more successful strategy with a probability depending on the pay-off differential. Profits of fundamental strategies are approximated by the relative deviation of the price from its fundamental value (arbitrage opportunities). Profits of chartism mainly depend on capital gains due to price changes.

Lux (1998) finds that his model produces chaotic motion for many parameter combinations. In a stylized way, the socio-economic dynamics evolve as follows. Large price increases during the beginning stages of a bubble convert pessimistic chartists and fundamentalists into optimistic chartists who profit from the up-trend. For some time, this heats up the bubble. However, once the bubble has infected a certain number of traders, the exhaustion of the pool of additional traders causes a slow-down of the price trend. With a relatively large deviation from the fundamental, there is a high potential gain from the fundamentalists strategy. Both tendencies gradually change the profit differential in favor of the fundamentalist strategy, fostering transitions from chartists to fundamentalists. As a result, the price trend eventually reverses. This, in turn, leads to an additional erosion of confidence among the remaining chartists.

This model has the potential to mimic the behavior of financial data quite closely. In Lux (1998), the distribution of returns exhibits fat tails. The setup of Lux (1997) produces volatility clustering and Lux and Marchesi (2000) are able to replicate unit roots in the asset price, fat tails for returns and volatility clustering.

2.2.2.8 Summary

The models of Frenkel, Cutler, Poterba and Summers, Day and Huang, Chiarella, and de Grauwe, Dewachter and Embrechts all assume that the group sizes of chartists and fundamentalists remain constant over time. However, this stands in contradiction to the findings of empirical studies. Survey studies like those of Taylor and Allen (1992) report that most of the agents are familiar with both technical and fundamental trading strategies. Only a minor fraction of the agents relies exclusively on one class of trading rules.

Since the pricing function of Frenkel, and Cutler, Poterba and Summers is linear the dynamics of their models appears less interesting. Without exogenous shocks, prices converge towards its long-run equilibrium. Day and Huang introduce a nonlinearity into their model by assuming that the demand of the fundamentalists is nonlinear. In the model of Chiarella, the demand of the chartists is nonlinear. Both approaches have the potential to generate interesting dynamics. For some parameter combinations the motion is chaotic. However, to be able to derive some analytic results, the behavior of the agents is drawn very simply. For instance, in Day and Huang chartists classify a market as a bull market if prices are above the long-run equilibrium. Thus, the pricing function becomes a one-dimensional nonlinear map. De Grauwe, Dewachter and Embrechts assume that fundamentalists are heterogeneous. This is an important point. If fundamentalists operate against each other, their influence is less strong. In addition, the expectation formation of the chartists is realistically described by a univariate time series model. Of course, such a model can only be analyzed with numerical means. Simulations show clear evidence of chaos.

The dynamics of the models of Frankel and Froot, Kirman, Brock and Hommes, and Lux are derived by the fact that agents switch between trading strategies. In the framework of Frankel and Froot, the portfolio managers aggregate the forecasts of chartists and fundamentalists with respect to the quality of their predictions. The model of Frankel and Froot tries to replicate the bubble path of the US dollar in the 1980s. As long as the influence of the chartists rises, the exchange rate is driven away from its fundamental. The weight of the chartists and the exchange rate move from period to period in the same direction. Such an explanation is highly stylized. Even during a bubble episode, the exchange rate path fluctuates in a complicated way and the agents switch between trading strategies. For Kirman, the switching between trading strategies heavily depends on the social interactions between the traders. For instance, traders may herd together on one strategy for some time. Kirman's model generates complex dynamics. However, the generated bubbles should only be interpreted as short-term deviations from fundamentals. Otherwise, the exchange rate path appears as unrealistic as the one of Frankel and Froot.

Brock and Hommes, and Lux suggest very sophisticated models. Their work currently marks the state of the art in the chartists-fundamentalists debate. In Brock and Hommes, predictors are selected with respect to past realized profits. Lux also considers social interactions between traders. A minor flaw of Lux' model is that the profits associated with technical and fundamental behavior are calculated differently. In the absence of shocks, both models can produce chaotic motion. Buffeted with noise, the price dynamics mimic the behavior of real financial data quite closely. Unfortunately, the dynamics of such models become increasingly difficult to replicate. From a scientific perspective, this is obviously not desirable. Future work should thus either report computational issues comprehensively or turn to less complex models.

2.2.3 Agent-Based Computation

2.2.3.1 Basic Elements

The agent-based computation approach may be defined as the study of economies represented as evolving systems of a large number of interacting agents. The idea of such a bottom up perspective is that global macro behavior is grounded in local agent interactions. Agent-based computation models have three key elements in common (compare the surveys of Tesfatsion 2001 or LeBaron 2000). First, agents are typically described as heterogeneous entities that determine their interactions with other agents and with their environment on the basis of internalized data and behavioral rules. These agents have a more elaborated cognitive structure than conventionally modeled economic agents. Second, the agents face a natural selection pressure, which results in the continual creation of new types of agents behavior. This evolutionary process does not take place on the population level but directly affects the characteristics of the agents behavior. Third, agent-based computation models are computer implemented virtual economic worlds that grow themselves over time. After the initial conditions are determined, all subsequent events in these virtual economies evolve by the action of the traders.

2.2.3.2 The Model of the Santa Fe Institute

To illustrate the philosophy of such models in more concrete terms, we briefly address the main features of its most prominent approach: the Santa Fe artificial stock market of Palmer et al. (1994), Arthur et al. (1997) and LeBaron et al. (1999).[11] They propose a theory of asset pricing based on heterogeneous traders. Each agent at any time possesses a multiplicity of forecast rules and uses those that are both best suited to the current state of the market and have recently proved most reliable. Agents learn by discovering which of their predictors prove best, and by developing new ones from time to time.

More specifically, each predictor is a condition-forecast rule that contains both a market condition that may at times be fulfilled by the current state of the market and a forecasting formula for the price of the next period. Each agent holds M individual predictors in mind and uses the most accurate of those that are active. Thus, each agent has the ability to recognize different sets of states of the market. For example, if a head and shoulders pattern has a high performance, a trader will not use it until he has really detected a head and shoulders pattern. If a head and shoulders pattern emerges, the forecasting part of the predictor triggers a trading position.

[11] Other contributions include Arifovic (1996), Youssefmir and Huberman (1997), and Farmer (2000).

The market price is driven by the trading activity of the agents. Once the market clearing price is revealed, the traders update the accuracy of the active predictors. In addition, the traders continually explore new forecasting rules. That is, from time to time they drop rules that perform badly, and create new ones. Note that learning takes place in two ways. On a rapid base as agents learn which of their predictors are accurate and worth acting upon, and which should be ignored, and at a slower time scale as non-performing predictors are discarded and new ones are invented.

The model generates very realistic price dynamics. Stylized facts like excess volatility, fat tails for returns, volatility clustering or high trading volumes are replicated. Another important result is that technical trading rules are not forced to be but endogenously persist in the market. For instance, technical analysis can emerge if trend following predictors are generated by chance in the population, and if random perturbations activate them and subsequently validate them.

2.2.3.3 Summary

Le Baron (2000), one of the Santa Fe Institute members, notes that the agent-based computation approach has both advantages and disadvantages. On the one hand, agents are allowed to explore a fairly wide range of possible forecasting rules. In addition, they have flexibility in using and ignoring different pieces of information. On the other hand, it is often difficult to pin down causalities acting inside the market. For instance, the model of the Santa Fe Institute investigates the interaction of 25 traders, each relying on a set of 100 predictors. Overall, 2,500 different predictors can be used at time t. As time proceeds, the traders constantly explore new rules.

Lawrenz and Westerhoff (2001) try to weaken 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. The agent-based computation model of Lawrenz and Westerhoff is able to replicate some of the most important stylized facts of the foreign exchange market. Since the traders only follow a set of 6 time-invariant trading rules, the model allows the development of an intuitive understanding of what is going on in the market.

2.2.4 Conclusions

The noise trader approach has sharpened the theoretical foundations of noise trading. Transactions of noise traders can shift prices away from fundamentals, because arbitrage is risky and therefore limited. Noise trading is not driven out of the market. Higher risk-taking may be rewarded by the market and in some situations it is optimal to use simple feedback rules. However, we find the underlying rationality assumption somewhat inconsistent. On the one hand, there exist unintelligent noise traders who do not learn that their behavior is erroneous. On the other hand, there are rational traders who possess full knowledge of the market.

Chartists-fundamentalists models study market fluctuations caused by the interactions of boundedly rational agents relying on a small set of popular trading rules. For instance, complicated dynamics may emerge due to nonlinear demand functions. We favor the view where the agents are actively able to choose between trading rules. Variations in the importance of technical and fundamental analysis lead to complex dynamics. Agent-based computation models extend the cognitive capabilities of the traders. Each agent keeps a record of the fitness of a large set of predictors. Bad predictors are regularly substituted against new ones. These models give a very detailed description of the behavior of the agents. Simulations of economies with large numbers of traders yield quite realistic asset price dynamics. The agent-based computation approach is very interesting, but one should note its drawbacks. Given only the information about a model available in a paper, it is almost impossible to replicate the dynamics. Such an approach can also be very time-consuming. Neely et al. (1997) report that it would take at least 81 days of computing time to duplicate their results on a 120 MHz Pentium. One has to hope that the authors (or referees) treat computational issues very carefully. Although much less complex in design, this also holds for some of the chartists-fundamentalists models.

Nevertheless, all three approaches have provided important new insights into the working of speculative markets. The competition of these co-evolving strands of research will hopefully stimulate mutual progress in the future.

2.3 Outline of the Program

In the next three chapters we present our own work. In chapter 3, a simple deterministic nonlinear exchange rate model is developed. Speculators repeatedly choose between technical and fundamental trading strategies to determine their investment positions. The choice of the rules depends on expected future performance possibilities. The interaction between the trading rules produces complex dynamics where the exchange rate is circling around its fundamental value without any tendency to converge. The model endogenously replicates the stylized facts of excess volatility, structural breaks in the level of the exchange rate and volatility clustering.

In chapter 4, the deterministic model is buffeted with dynamic noise. Due to this extension the model is able to mimic various stylized facts. Simulations produce a high variability of the exchange rates, fat tails for returns and weak evidence of mean reversion. Within this framework, the impact of central bank interventions like "leaning against the wind" on the exchange rate variability is analyzed. This is one of the first contributions which uses a chartists-fundamentalists model as a laboratory to investigate, evaluate and improve the effectiveness of intervention operations.

Chapter 5 explores the phenomenon of lasting deviations of the exchange rate from its fundamental value. In chapter 3 and 4, the strong assumption is made that the agents are able to determine the fundamental value of the exchange rate. Chapter 5 adds to the literature in the way that psychological factors which influence the perception of the fundamental exchange rate are modeled explicitly. Simulations give rise to bubbles but simultaneously display quite realistic exchange rate dynamics.

Note that chapters 3-5 appear as independent papers including abstract, keywords and JEL classification. This brings both advantages and disadvantages. Each chapter is, of course, self-explaining, but we partially repeat ourselves. Although the models of chapter 3-5 are different in several aspects, they are built on the same basic principles. Hence, some of our arguments necessarily reappear in a similar way. However, we hope that the clarity connected with our style of presentation outweighs this disadvantage.

3 Speculative Behavior and Exchange Rate Dynamics

Abstract

This chapter* deals with speculative trading. Guided by empirical observations, a deterministic nonlinear exchange rate model is developed in which the speculators repeatedly choose between technical and fundamental trading strategies to determine their investment positions. The interaction between the trading rules produces complex – sometimes even chaotic – exchange rate dynamics. The model endogenously replicates the stylized facts of excess volatility, structural breaks in the level of the exchange rate, and volatility clustering. Overall, our results are robust for different functional specifications and parameter settings.

Keywords

exchange rate theory, nonlinear dynamics and chaos,
technical and fundamental trading rules

JEL Classification

C63, F31, G14

* Presented at the 1st CeNDEF Workshop on Economic Dynamics, Amsterdam, January 2000, and at the ZEI Summer School on International Macroeconomics, Bad Honnef, August 1999. I would like to thank the participants for their helpful discussions, especially Cees Diks, Doyne Farmer, Paul de Grauwe, Jürgen von Hagen, Cars Hommes, Alan Kirman and Thomas Lux.

3.1 Introduction

The foreign exchange market can be regarded as a virtually 24-hour global market to which all major trading centers around the world are connected by electronic links. Since the development of real-time information systems and the decline in transaction costs following the liberalization of the capital markets in the mid 1980s, the daily turnover has increased sharply. More and more, the trading volume reflects very short-term transactions indicating a highly speculative component, whereas international trade transactions account for merely one percent of the total (BIS 1999).

When determining their speculative investment positions the market participants rely on both technical and fundamental trading rules. Technical analysis is a trading method that attempts to identify trends and reversals of trends by inferring future price movements from those of the recent past.[1] By contrast, fundamental concepts look at the underlying reasons behind that action. According to Taylor and Allen (1992), most foreign exchange dealers place at least some weight on technical analysis. Especially for short-run predictions, technical trading rules are regarded as very useful. Furthermore, it is quite common for professional traders to use both trading methods. The decision to choose one of the rules depends on the expected chances of success.

This chapter is based on such observations. Our aim is to investigate how speculative trading affects exchange rate fluctuations. An commonly heard conjecture is that speculation causes excess volatility. Moreover, speculation may also be related to other stylized facts such as volatility clustering. In order to develop a realistic framework, we use empirical evidence to model the behavior of the agents. Our starting point is that traders rely on simple rules to determine their investment positions.[2] To develop a better understanding of the dynamics, we study the interaction between these rules within a deterministic setting.

---

[1] For a short introduction to technical analysis see Neely (1997). A comprehensive discussion is found in Murphy (1999), which is often referred to as the "bible" of technical analysis.

[2] In a broader sense, this chapter enriches the noise trader approach (Shleifer and Summers 1990) with details from both the market microstructure literature and psychological evidence. Models which also study the interaction between chartists and fundamentalists include Frankel and Froot (1986), de Grauwe et al. (1993), Kirman (1993), Lux (1997), Brock and Hommes (1998), and Farmer (2000).

Our main results are: simulations produce equilibrium exchange rates which erratically fluctuate around some fundamental value without any apparent tendency to converge. The trading signals needed to keep this process going are generated by the agents themselves. Thus, we argue that exchange rate dynamics are at least partially caused by an endogenous nonlinear law of motion. Besides replicating the stylized fact of excess volatility, the dynamics display endogenous changes in the level or in the volatility of the exchange rate. In addition, the time series are explored with the methods of nonlinear time series analysis; strong features of chaos such as positive Lyapunov exponents or low dimensional attractors, are detected. These findings are robust for different functional specifications and parameter settings.

This chapter is organized as follows. Section 3.2 suggests a general framework for nonlinear exchange rate dynamics. In sections 3.3 and 3.4, we specify and discuss two examples of the framework. Section 3.5 concludes the chapter. Some technical notes on nonlinear time series analysis are presented in the appendix.

3.2 Nonlinear Exchange Rate Dynamics

Our focus is on speculative trading. We consider traders who are familiar with both technical and fundamental analysis. The agents determine their positions by applying one specific trading rule. Selection of the rules depends on expected performance possibilities. The decision to choose a certain trading rule is repeatedly made every period before the trading starts.

As already mentioned, the agents have the choice between technical and fundamental trading rules. Hence, we call them chartists or fundamentalists. Technical trading rules rely on past movements of the exchange rate as an indicator of market sentiment and extrapolate these into the future, thus adding a positive feedback to the dynamics. Fundamental trading rules are designed to exploit the differences between the spot price and the fundamental value of a currency. Since fundamentalists trade on a reduction of the mispricing they add a negative feedback to the dynamics (mean reversion).

The decision to select a particular trading rule depends on expected future profit opportunities, which the agents try to derive out of the condition (the mood) of the market. One often observes that fundamentalism, compared to chartism, becomes more popular the wider the spot rate deviates from its perceived fundamental value. In the language of the traders, the market becomes oversold or overbought (Murphy 1999). In such a situation, the agents believe that the chance that the exchange rate returns to its fundamental value increases as the misalignment rises.[3]

For simplicity, other demand, such as firms transactions relating to international trade and risk management or central bank interventions, is excluded. This might be justified by the fact that these transactions, in contrast to the speculative trading positions, are small in absolute magnitude (BIS 1999). Every period, the exchange rate is determined via the market clearing condition.

Note that the selection of the rules introduces nonlinearity into the dynamics. The chartists are most influential if the exchange rate is near its perceived fundamental value. Since the behavior of technical traders is trend extrapolating, the exchange rate is typically driven away from its fundamental value. However, the higher the misalignment of the exchange rate, the more the market impact of the fundamentalists increases. Transactions of this group lead to a mean reversion until the chartists again reign over the market.

To concentrate on the interaction of the trading rules we study a simple deterministic setting. For further analysis, the trading rules and the selection process have to be formalized. For each of the three equations we introduce two functional specifications. Out of the eight versions, we discuss two of them as examples.[4] The aim of this exercise is twofold. First, to develop a better understanding of exchange rate fluctuations, and second, to check the robustness of the dynamics.

---

[3] Alternatively, the agents may use past realized profits as an indicator for successful strategies. Lawrenz and Westerhoff (2001) discuss such behavior in a related setting.

[4] Together with some notes on numerical tools for detecting chaos, all eight versions are briefly compared in the appendix. The versions are labeled according to the first letters of their components, that is P or G for a popular or a general technical trading rule, R or A for regressive or anchor expectations, and Q or S for a quadratic or a square root weighting scheme.

3.3 The PRQ-Version

3.3.1 Specification

One of the most popular technical trading rules is the double crossover method (Murphy 1999), in which a buy (sell) signal is given when a short-term moving average of the past exchange rates crosses a long-term moving average of the past exchange rates from below (above). This rule may be formalized as

$$d_t^C = \alpha^C((LogS_{t-1} - LogS_{t-2}) - 0.5((LogS_{t-1} - LogS_{t-2}) + (LogS_{t-2} - LogS_{t-3})))$$
$$= \alpha^C(0.5LogS_{t-1} - LogS_{t-2} + 0.5LogS_{t-3}), \quad (1)$$

where the first term is a one-period change reflecting the short-term moving average, and the second term is a similarly constructed long-term moving average.[5] The demand is calibrated from the reaction coefficient α^C. Note that by (1) chartists place a market order today in response to past price changes, i.e. 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).

Fundamental analysis is built on the premise that the exchange rate will converge towards its equilibrium value S^F in the future. For simplicity, we assume that the agents correctly perceive the (constant) fundamental exchange rate. Fundamental trading rules differ in the way in which they deal with this adjustment process. According to Takagi (1991), the agents often form regressive expectations like

$$E_t[S_{t+1}] = \gamma S^F + (1-\gamma)S_{t-1}, \quad (2)$$

where γ represents the expected adjustment speed of the exchange rate towards its fundamental value. Since expectations have to be formed before trading starts, the last available data is from period t-1. Then, the demand of the fundamentalists may be formulated as

$$d_t^F = \alpha^F(E_t[S_{t+1}] - S_t)/S_t = \alpha^F(\gamma S^F + (1-\gamma)S_{t-1} - S_t)/S_t. \quad (3)$$

[5] This study abstracts from qualitative analysis of past trends, like the famous head and shoulders formation. Restricting the analysis to trading strategies like (1) need not necessarily be a disadvantage because even simple rules give rise to complex exchange rate movements. The incorporation of more complex rules should, however, intensify the dynamics.

Fundamental trading rules deliver a buy (sell) signal if the expected future exchange rate is above (below) the spot rate. The demand is calibrated according to α^F and depends on the relative distance between $E_t[S_{t+1}]$ and S_t.[6]

The decision to select a trading rule depends on expected future performance possibilities and has to be made before the trading starts. The weight of the chartists is defined as

$$m_t = \frac{1}{1+\beta^1+\beta^2((E_t[S_{t+1}]-S_{t-1})/S_{t-1})^2}, \tag{4}$$

whereas that of the fundamentalists is ($1-m_t$). The coefficient β^1 reflects the basic proportion of agents who are always fundamentalists. If, for example, β^1 is 0.25, then 20 percent of agents are permanently fundamentalists whatever the situation of the market. Nevertheless, most traders adjust their trading strategies with respect to the condition of the market. If the exchange rate is close to its expected future value, fundamental analysis provides no clear trading signals. The agents then tend to prefer technical analysis to identify the future direction of the market. However, the more the exchange rate deviates from its expected future value, the more convincing fundamental analysis becomes. The coefficient β^2 indicates the popularity of fundamentalism. The weighting function looks like a bell-shaped curve with respect to the relative distance between $E_t[S_{t+1}]$ and S_{t-1}.

The market clearing condition is given as the sum over all trading positions

$$m_t d_t^C + (1-m_t) d_t^F = 0. \tag{5}$$

Combining (1) - (5) and solving for the exchange rate yields

$$S_t = \frac{\gamma S^F + (1-\gamma)S_{t-1}}{1-\dfrac{\alpha^C(0.5LogS_{t-1}-LogS_{t-2}+0.5LogS_{t-3})}{\alpha^F(\beta^1+\beta^2((\gamma S^F+(1-\gamma)S_{t-1}-S_{t-1})/S_{t-1})^2)}}, \tag{6}$$

which is a three-dimensional nonlinear deterministic difference equation. Since (6) precludes closed analysis, we simulate the dynamics to demonstrate that the underlying structure gives rise to complex exchange rate behavior, as is observed empirically.

[6] Note that only two groups act in the market. Since chartists trade on the basis of past price movements, their demand is perfectly inelastic. Thus, we have assigned the job of market maker to the fundamentalists, who absorb the demand of the chartists by setting the new exchange rate.

3.3.2 Simulations

We will start with some comments on the underlying parameter setting. Unfortunately, not all coefficients are empirically observable. To obtain a first base run the following assumptions have been made. The coefficients α^C and α^F are equal, the expected adjustment speed of the exchange rate towards its fundamental value is 20 percent and S^F=100. Survey studies (Taylor and Allen 1992) report that 5 to 15 percent of the agents rely solely on fundamental analysis. Hence, β^1=0.125 seems to be a good choice. The extent of the exchange rate volatility is then calibrated via the popularity of fundamental analysis. For β^2=320,000, the dynamics evolve quite realistically.

The top chart of figure 1 displays a typical example of the exchange rate dynamics for 500 periods. After an initial shock, the exchange rate circles in a complex fashion around its fundamental value. The bottom chart shows the corresponding weight of chartists. The agents often switch between technical and fundamental analysis. At no time is one of the rules ever driven out of the market.

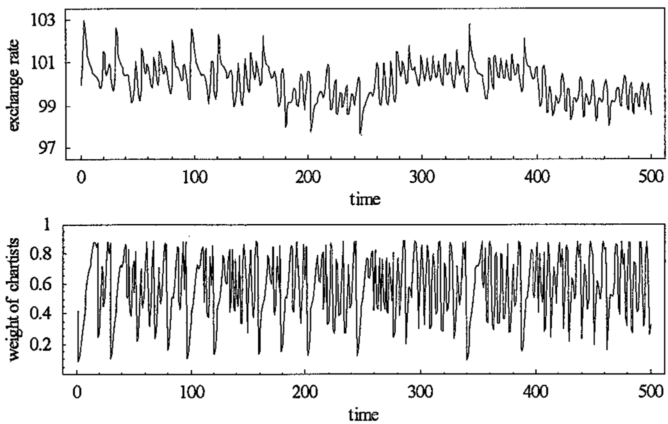

Figure 1: Exchange Rate and Weight of Chartists (PRQ-Version). S^F=100, S_1=100, S_2=101, S_3=103, α^C=1, α^F=1, γ=0.2, β^1=0.125, β^2=320,000.

The dynamics might be explained as follows. Technical trading rules always produce some kind of buy or sell signal and may, on the basis of a feedback process, induce a self-reinforcing run. 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 its expected future value, the stronger the influence of the fundamentalists, until eventually their increasing net position triggers a mean reversion. However, this already indicates a new signal for the chartists and leads directly to the next momentum. The exchange rate is repelled by its fundamental value because chartism dominates the market in this region. Heavy outliers occur when the chartists have a clear trading signal and the influence of the fundamentalists is low.

It is worth noting that our simple model already suffices to produce the high volatility of exchange rates that is observed empirically. A single disturbance is enough to trigger an exchange rate movement that circles around a fundamental value without converging. The volatility of the foreign exchange market need not be caused by exogenous shocks; it might be explained at least partially by an endogenous nonlinear law of motion. The trading signals needed to keep the process going are generated by the agents themselves.

This is exactly what Black (1986) has called noise trading. Black concludes that noise in the sense of a large number of small events is essential to the existence of liquid markets. He argues 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 different information. In our model, noise is permanently produced by the agents themselves. Even when there is no new information at all, trading volume and volatility will be high. Noise trading is trading on noise as if it were information.

According to Heiner (1983), such behavior may not necessarily be irrational. Heiner argues that the limits to maximizing in an uncertain environment are the origin of a rule-governed behavior. For example, for every agent the specific complexity of the foreign exchange market leads to a gap between his competence to make an optimizing decision and the actual difficulty involved with this decision. Thus, agents cannot do much better in the presence of complex dynamics than to follow some adaptive scheme of behavior.

Is this also true within our framework? To answer this question we have designed the following experiment. Suppose that one agent always relies on technical analysis as specified by (1). The technical trading rule delivers a right trading signal if it correctly predicts the change between S_{t+1} and S_t (else the signal is wrong). For example, the trading rule delivers a right signal if it generates a buy signal and the exchange rate rises. Figure 2 displays the outcome for trading periods 401-500. Within this sample, technical analysis produces 51 right decisions. Hence, the agent is more often right than wrong; his behavior should not be called irrational. The profitability of technical analysis has also been demonstrated empirically (Brock et al. 1992).

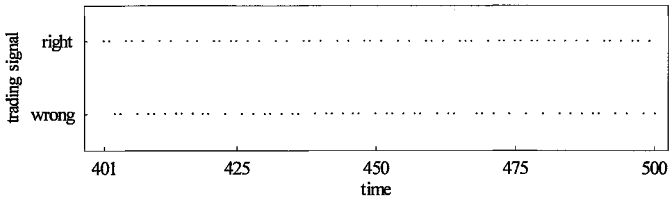

Figure 2: Technical Trading Signals. S^F=100, S_1=100, S_2=101, S_3=103, α^C=1, α^F=1, γ=0.2, β^1=0.125, β^2=320,000.

3.3.3 Quantifying the Dynamics

Nonlinear dynamic systems have the potential to produce chaotic motion. Next, we want to quantify the dynamics of our model. Although no commonly agreed definition of chaos exists, three important aspects regularly emerge (for a short introduction see Baumol and Benhabib (1989), else compare the references in the appendix):

- first, the trajectory of a deterministic process should be highly irregular. At least some of the standard tests of randomness cannot distinguish between chaotic patterns of change and truly random behavior,
- second, the time path is sensitive to a microscopic change in the value of the initial conditions (SIC). This means that for a slightly different choice of an initial value, the trajectory of the system would initially remain very close to the original trajectory in phase space, but eventually diverge from it,
- third, the time path may never return to any point it has previously crossed, but displays an oscillator pattern in a bounded region. In a phase space representation, the dynamics display some order (structure). For instance, a strange attractor may emerge.

Before we describe the generated time series more formally with methods of nonlinear time series analysis, we want to develop a better intuitive understanding of what sensitivity to initial conditions (SIC) means. Figure 3 compares two simulation runs with nearly identical sets of initial conditions and parameters. The only difference is that the solid line is computed with S_3=102 and the dashed line with S_3=102.0001. Surprisingly, after about 50 periods the time series start to diverge. In the bottom chart the difference between the two time series is plotted. After some iterations, the difference grows to the same order of magnitude as the usual fluctuations. A small change in the initial conditions changes the whole future path of the exchange rate in a quite dramatic way.

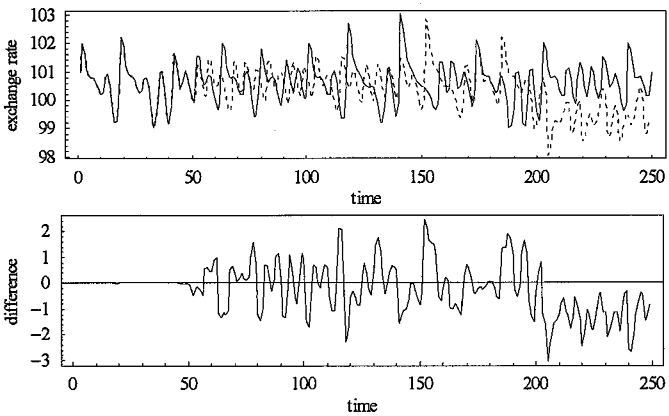

Figure 3: Example of SIC (Initial Value). Solid line: S^F=100, S_1=100, S_2=101, S_3=102, α^C=1, α^F=1, γ=0.2, β^1=0.125, β^2=320,000. Dashed line: the same but S_3=102.0001.

Another manifestation of SIC is the way in which a small change in the parameters of the model affects the time path of the exchange rate. Again, figure 4 compares two simulation runs with almost identical initial values and parameters. Now the only difference is that the solid line is generated with β^1=0.125 and the dashed line with β^1=0.125025. As visible from figure 4, the outcome is very much like the outcome displayed in figure 3. After around 50 periods, the two time series evolve completely differently. This explains further why forecasting nonlinear time series is a difficult task. Suppose that a trader has an almost perfect knowledge about the underlying model.

Even a small lack of precision may prevent the trader from successfully exploiting his knowledge. Reasonable predictions are only possible in the short run.

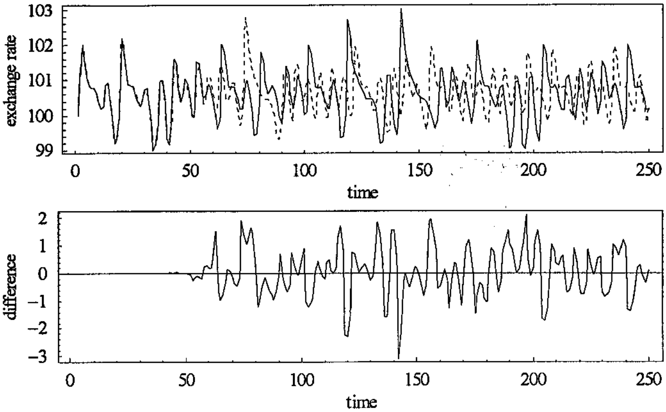

Figure 4: Example of SIC (Parmeter). Solid line: S^F=100, S_1=100, S_2=101, S_3=101.5, α^C=1, α^F=1, γ=0.2, β^1=0.125, β^2=320,000. Dashed line: the same but β^1=0.125025.

Numerical tests for chaos (mainly) consist of computing the largest Lyapunov exponent. The Lyapunov exponent measures how sensitively a trajectory reacts to a change in initial conditions. If the Lyapunov exponent is positive, two originally close orbits start to diverge in phase space. A (bounded) time series is said to be chaotic if the largest Lyapunov exponent is positive. The exponent is negative for fixed points or periodic data, zero for quasi-periodic data, and infinite for uncorrelated random data.

Figures 5-8 show Lyapunov exponents for various parameter settings. Figure 5 displays Lyapunov exponents for various values of β^1 together with a bifurcation diagram. A bifurcation diagram is a powerful graphical tool to visualize the dynamic properties of a system. It is constructed by increasing the coefficient β^1 in 1,000 steps from 0.01 to 1.00. Each time, we calculate 1,050 data points. To exclude possible transient phases of the trajectory, only the last 50 observations are plotted. Note that the motion is not chaotic for all β^1. The dynamics may also be periodic or quasi-periodic. If β^1 is higher than around 0.62, the solution becomes a stable fixed point.

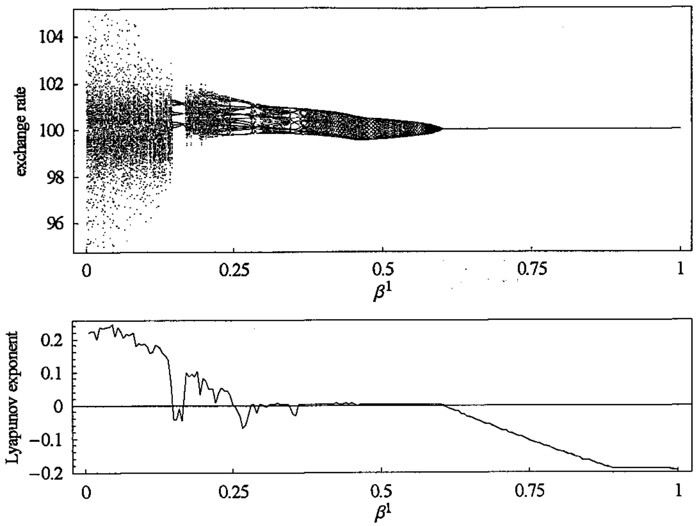

Figure 5: Lyapunov Exponents and Bifurcation Diagram for β^1. S^F=100, S_1=100, S_2=101, S_3=101.5, α^C=1, α^F=1, γ=0.2, β^2=320,000. Top: β^1 rises in 1,000 steps as marked on the axis, $S_{1,001}$-$S_{1,050}$ are plotted. Bottom: Lyapunov exponents are calculated for $S_{1,001}$-$S_{2,000}$.

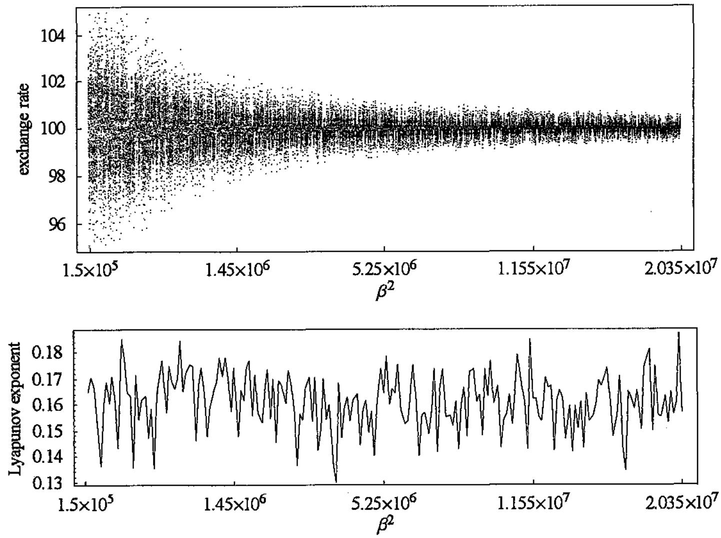

Figure 6: Lyapunov Exponents and Bifurcation Diagram for β^2. S^F=100, S_1=100, S_2=101, S_3=101.5, α^C=1, α^F=1, γ=0.2, β^1=0.125. Top: β^2 rises in 1,000 steps as marked on the axis, $S_{1,001}$-$S_{1,050}$ are plotted. Bottom: Lyapunov exponents are calculated for $S_{1,001}$-$S_{2,000}$.

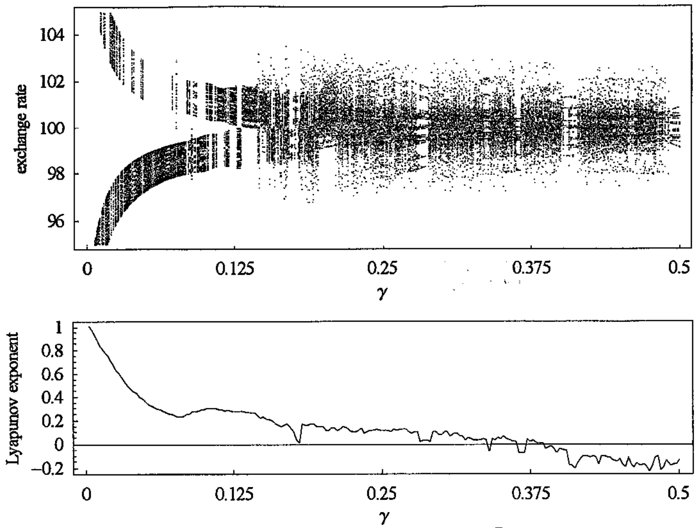

Figure 7: Lyapunov Exponents and Bifurcation Diagram for γ. S^F=100, S_1=100, S_2=101, S_3=101.5, α^C=1, α^F=1, β^1=0.125, β^2=320,000. Top: γ rises in 1,000 steps as marked on the axis, $S_{1,001}$-$S_{1,050}$ are plotted. Bottom: Lyapunov exponents are calculated for $S_{1,001}$-$S_{2,000}$.

Figure 6 presents the bifurcation diagram and the Lyapunov exponents for rising β^2. Remember that an increase of β^2 stands for an increase in the popularity of fundamental analysis. The influence of the fundamentalists increases from the left to the right. Periods of high volatility occur when fundamentalism is less popular. All Lyapunov exponents are positive independently of β^2. Although higher β^2 reduce the volatility, the qualitative nature of the dynamics remains stable. This is not the same if γ is varied (figure 7). For lower γ, the motion has a tendency to become more irregular.

Figure 8 contains Lyapunov exponents in the β^1/γ-plane; the black (white) dots on the right represent positive (negative) Lyapunov exponents. Roughly speaking, the model produces chaotic motion for low values of β^1 and γ. If the fraction of fundamentalists increases the fluctuations are stabilized. The evidence for chaos also vanishes if the agents expect a faster adjustment process. In this case, the exchange rate is more strongly attracted towards its fundamental value. However, for reasonable parameter settings the dynamics are likely to exhibit chaos.

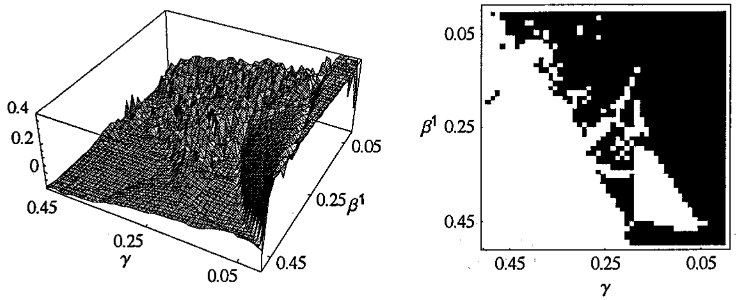

Figure 8: Lyapunov Exponents in the β^1/γ - Plane. S^F=100, S_1=100, S_2=101, S_3=101.5, α^C=1, α^F=1, β^2=320,000. β^1 and γ rise in 100 steps as marked on the axis. The Lyapunov exponents are calculated for $S_{1,001}$ - $S_{2,000}$. Right: black dots indicate positive Lyapunov exponents.

Figure 9 displays the dynamics in phase space, that is S_t is plotted against S_{t-1}. The same is carried out for the logarithm of price changes (the returns). The top of the figure contains the whole attractor, the bottom shows a blow-up. Although the dynamics in the time domain seem rather random, a distinct structure builds up in the phase space. Due to SIC, no meaningful precise forecasting is possible in the long run. However, some long-term properties can be identified, such as the structure in the phase space.

The correlation dimension describes the complexity of the attractor. The higher the dimension, the more complex the structure. A correlation dimension greater than about five implies essentially random data. Since our model is a three-dimensional difference equation system, the upper bound of the correlation dimension is three.

For our base line simulation, displayed in figure 1, the Lyapunov exponent is estimated as 0.142 and the correlation dimension as 1.62. Thus, the time series exhibits strong evidence of chaos.

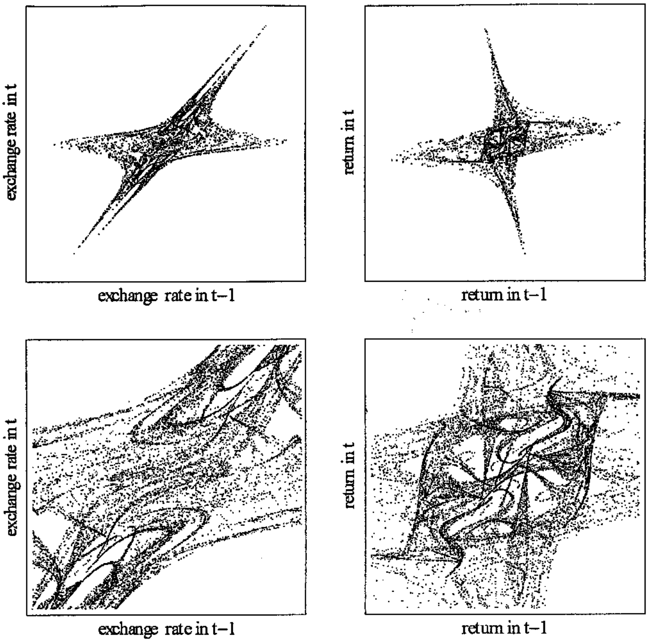

Figure 9: Exchange Rates and Returns in Phase Space (PRQ-Version). S^F=100, S_1=100, S_2=101, S_3=103, α^C=1, α^F=1, γ=0.2, β^1=0.125, β^2=320,000. Top: data from t=1,000-11,000. Bottom: a blow-up of the attractor from t=1,000-31,000.

3.3.4 Some Nonlinear Phenomena

This section aims at relating typical nonlinear phenomena to well-known stylized facts of exchange rate dynamics. The first group of phenomena is concerned with changes in the level of the exchange rate, the second group with changes in the volatility of the exchange rate.

Sudden changes in the level of the exchange rate are a commonly observed price pattern. Such breaks may be triggered by fundamental shocks. But, as reported by Goodhart (1988), large price movements unrelated to any news are also apparent. Regime breaks may also be explained by the peso problem. If the agents expect a fundamental shift in the economy, the exchange rate reacts immediately. Moreover, when the expectations prove to be wrong, the exchange rate moves back to its previous level. In such a case, the level of the exchange rate has changed twice although no fundamental shock has occurred.

Structural breaks can easily be modeled exogenously within our model. The top chart of Figure 10 shows a simulation where a single fundamental shock occurs in period 300, raising the fundamental value of the exchange rate from S^F=100 to S^F=104. In the first 300 periods, the exchange rate fluctuates as usual around its fundamental value. Afterwards the whole dynamics are raised to the new equilibrium. In the bottom chart, the fundamental exchange rate between periods 200 and 400 is S^F=96, otherwise it is S^F=100. If one interprets this interlude as a temporary misperception of the fundamental exchange rate, then the bottom chart is reminiscent of the peso problem.

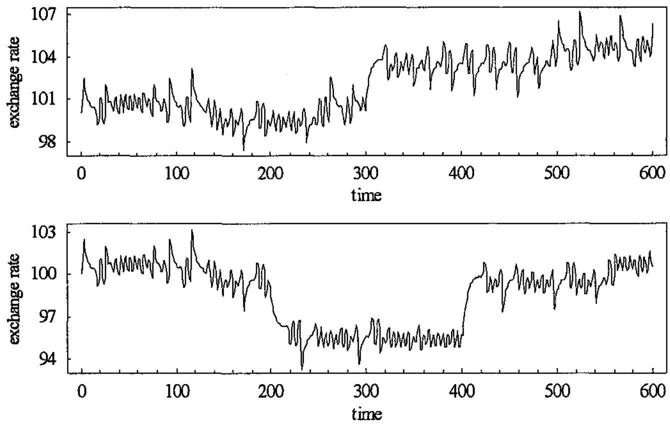

Figure 10: Changes in the Level of the Exchange Rate. S_1=100, S_2=101, S_3=103, α^C=1, α^F=1, γ=0.2, β^1=0.125, β^2=320,000. Top: for $t>300$ S^F=104, else S^F=100. Bottom: for $200<t<400$ S^F=96, else S^F=100.

A certain characteristica observable in nonlinear dynamical systems allows us to explain such a pattern with fewer exogenous interventions. In the case of coexisting attractors, one has the puzzling feature that the repetition of a simulation with the same parameters may yield a qualitatively different result. On which attractor the trajectory will settle depends on the initial conditions. Figure 11 shows such an example. The only difference between the simulation setup in the top and the middle charts is that S_3=101.2 in the top chart and S_3=99.6 in the middle chart. However, the outcome is quite distinct. After some iterations, one orbit evolves to a limit cycle above the fundamental value of the currency; the other lies below it.[7] In the bottom section of the figure, there seem to be fundamental breaks in the time series. For instance, between t=167 and t=333 the exchange rate fluctuates around a lower level than at other times. But no fundamental shock has occurred. Only in two trading periods is the foreign exchange market hit by exogenous noise, that is we have set S_{167}=99.5 and S_{333}=101.2. This is already enough for the trajectory to change its attractor. Apparently, structural breaks in a time series may be nothing more than jumping between different attractors.

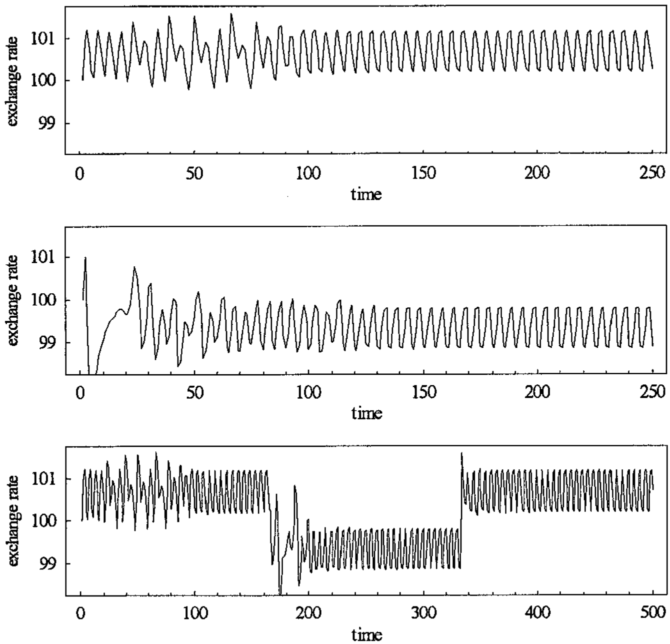

Figure 11: Coexistence of Attractors. Top: S^F=100, S_1=100, S_2=101, S_3=101.2, α^C=1, α^F=1, γ=0.2, β^1=0.167, β^2=320,000. Middle: the same but S_3=99.6. Bottom: the same but S_{167}=99.5 and S_{333}=101.2.

Any single time series can only present one possible attractor. The region in a phase space leading to a given attractor is called its basin of attraction. The basins of two

[7] For further analysis, a simple example of coexisting attractors (two limit cycles) is chosen. Of course, more complicated combinations of attractors exist. For β^1=0.375, two quasi-periodic attractors coexist and for β^1=0.1475, one finds a limit cycle combined with a chaotic attractor.

coexisting attractors can be interwoven in a very intricate way. The right-hand section of figure 12 shows the basins of attraction for the above setup in the S_{t-1}/S_{t-2}-plane. The initial conditions S_2 and S_3 are varied as marked on the axis. The limit cycle located below the fundamental value is represented by black dots, while the other is shown by white dots. The resolution of the plot is 400 times 400. The left section of figure 12 shows the time the trajectory needs to reach its attractor. The speed of adjustment is indicated by different levels of gray. The lighter the shade of gray, the faster the convergence. Although one observes a somehow symmetric picture for the basins of attraction, the speed of adjustment is more or less independent from where the orbit starts. After around 1,000 periods, all trajectories have settled down on their attractor.

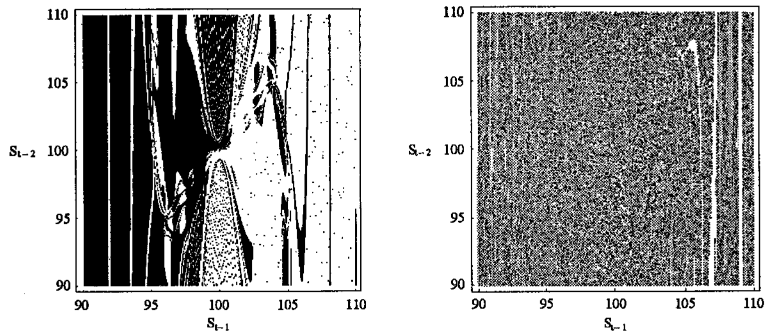

Figure 12: Basin of Attraction (PRQ-Version). $S^F=S_1=100$, $\alpha^C=\alpha^F=1$, $\gamma=0.2$, $\beta^1=0.167$, $\beta^2=320{,}000$, S_2 and S_3 as marked on the axis. Right: black (white) dots represent an attractor below (above) the fundamental value. Left: duration of transient behavior, lighter gray levels indicate faster convergence.

So far it is necessary to add at least some noise to mimic structural breaks. But the model has the potential to generate such phenomena completely endogenously. Figure 13 provides an example. In the top section, the exchange rate is plotted for 30,000 periods, while the bottom section shows the time series for t=4,100-4,500. For a long time period, the exchange rate fluctuates within a band between 100 and 104. Suddenly, without any apparent reason, the band shifts down for some time. Again, out of the blue, the dynamics switch back to its former region and stay there for a long time until a similar pattern repeats itself. Note that the fluctuations are solely driven endogenously; there is no noise added.

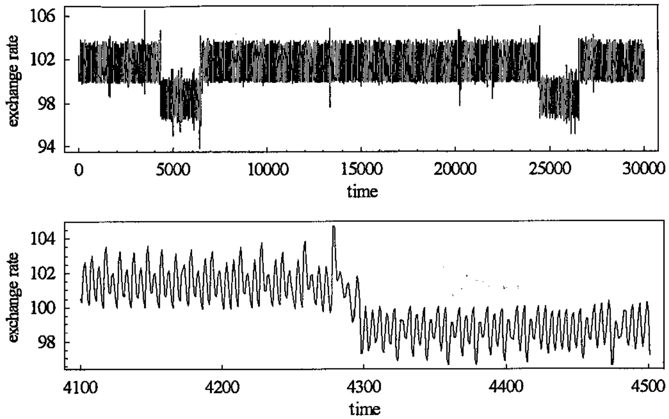

Figure 13: Endogenous Changes in the Level of the Exchange Rate. S^F=100, S_1=100, S_2=101, S_3=102.5 α^C=1, α^F=1, γ=0.2, β^1=0.1125, β^2=320,000.

Besides changes in the level of the exchange rate, one also observes variations in the volatility of the exchange rate. Volatility clustering describes the phenomenon in which periods of low volatility alternate with periods of high volatility (Mandelbrot 1963). One way to replicate such a pattern is, of course, to exogenously adjust the parameter of the model. Figure 14 displays two examples. The following shocks are assumed. In the top section for $2{,}000<t<4{,}000$ β^2=70,000, otherwise β^2=800,000, and in the bottom section for $200<t<400$ β^2=80,000, otherwise β^2=800,000. Hence, one source of volatility clustering might be changes in the popularity of fundamental analysis. Periods of high volatility coincide with periods in which the popularity of fundamentalism is low (and vice versa). Especially in periods of high uncertainty about the fundamental condition of a currency, fundamentalism is not very popular. In such a situation, a central authority may have the chance to reduce volatility by creating a greater consensus about the fundamental value of the currency, for example, by providing better public information. But β^2 may also autonomously switch its value. Sometimes, the traders tend to herd together. For instance, if a famous guru supports a certain class of trading strategies, then volatility clustering is unrelated to fundamental reasons.

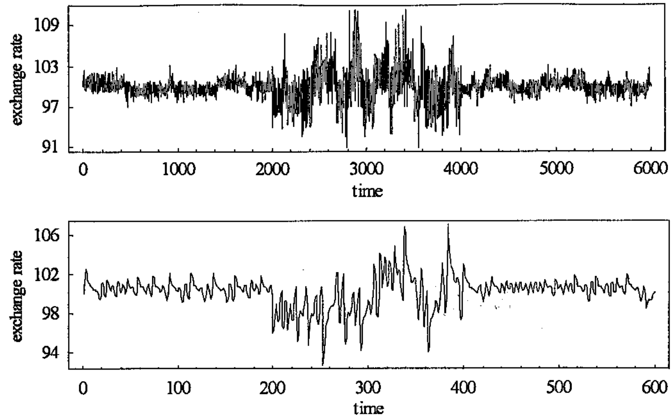

Figure 14: Volatility Clustering. S^F=100, S_1=100, S_2=101, S_3=102.5, α^C=1, α^F=1, γ=0.2, β^1=0.125. Top: for 2,000<t<4,000 β^2=70,000, otherwise β^2=800,000. Bottom: for 200<t<400 β^2=80,000, otherwise β^2=800,000.

Some nonlinear phenomena are also suggestive for volatility clustering: transient behavior and on-off intermittency. By definition, transient behavior disappears after some time. Therefore at first sight it might neither seem relevant nor interesting. But both impressions are wrong. Starting with some initial conditions, one has to wait some time until the trajectory has settled down on the attractor. During this transient phase the motion may have completely different properties than on the attractor itself. The transient time can be extremely short (for instance, for a stable fixed point). But in certain cases, transients can last a long time. The transient can be nontrivial even if the attractor itself is simple. Figure 15 contains a simulation run in which the attractor is a limit cycle. Before the attractor is reached, the dynamics are highly irregular. The bottom part of the figure shows a transient which lasts over 500 periods. During the transient phase the dynamics are comparatively volatile. Afterwards, the fluctuations are less pronounced. Note that a single disturbance can be enough for the system to switch back to turbulent motion. In period t=1,204, the exchange rate is set to 102.6555. The trajectory needs around 500 periods to approach its attractor. Hence, volatility clustering might be caused by temporary shocks which trigger complex transient dynamics in an otherwise stable environment.

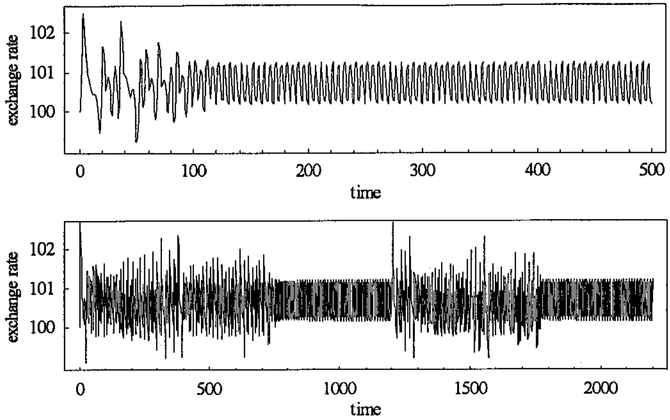

Figure 15: Transient Behavior. Top: S^F=100, S_1=100, S_2=101, S_3=102.5, α^C=1, α^F=1, γ=0.2, β^1=0.15875, β^2=320,000. Bottom: the same but S_3=103 and $S_{1,204}$=102.6555.

On-off intermittency means that the dynamics alternate between laminar and turbulent motion in an irregular fashion. The chaotic phases can be long or they can look like short bursts. Figure 16 shows an example for three different time periods. For a certain parameter combination the behavior of the model switches back and forth between two qualitatively different kinds of motion, even though all the parameters remain constant and no external noise is present. The switching appears to occur randomly. Both the duration and the frequency of bursts of chaotic behavior are unsystematic. On-off intermittency is an endogenous source for volatility clustering.

To sum up, the PRQ-version is able to replicate some important stylized facts of exchange rate dynamics: excess volatility, changes in the level of the exchange rate, and volatility clustering. On the one hand, these might be caused by external stochastic disturbances. On the other hand, nonlinear dynamic systems have the power to mimic such fluctuations completely endogenously. Finally, figure 17 presents a simulation run over 1,000,000 periods. The dynamics are entirely deterministic and contain both shifts in the level of the exchange rate and volatility clustering.

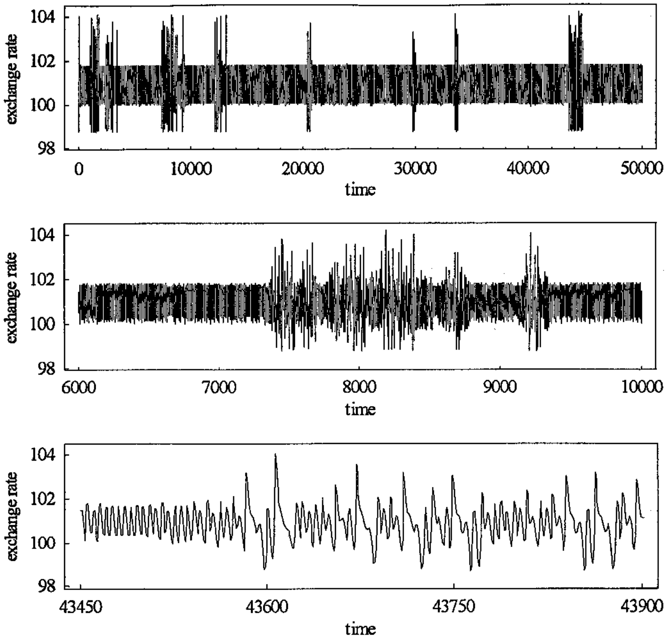

Figure 16: On-off Intermittency. S^F=S_1=100, S_2=101, S_3=103, α^C=α^F=1, γ=0.2, β^1=0.145, β^2=170,000.

3.4 The GAS-Version

3.4.1 Specification

Instead of a specific popular technical trading rule we now use a very general description of what technical traders do (Murphy 1999). Typically, chartists buy foreign currency if the exchange rate rises (and vice versa)

$$d_t^C = \alpha^C (0.6(LogS_{t-1} - LogS_{t-2}) + 0.4(LogS_{t-2} - LogS_{t-3})). \qquad (7)$$

This strategy belongs to the class of moving average rules, where the traders assign a greater weight to more recent exchange rate changes (thus 0.6 to 0.4). The demand for foreign currency is positive if the sum of the extrapolating terms is positive.

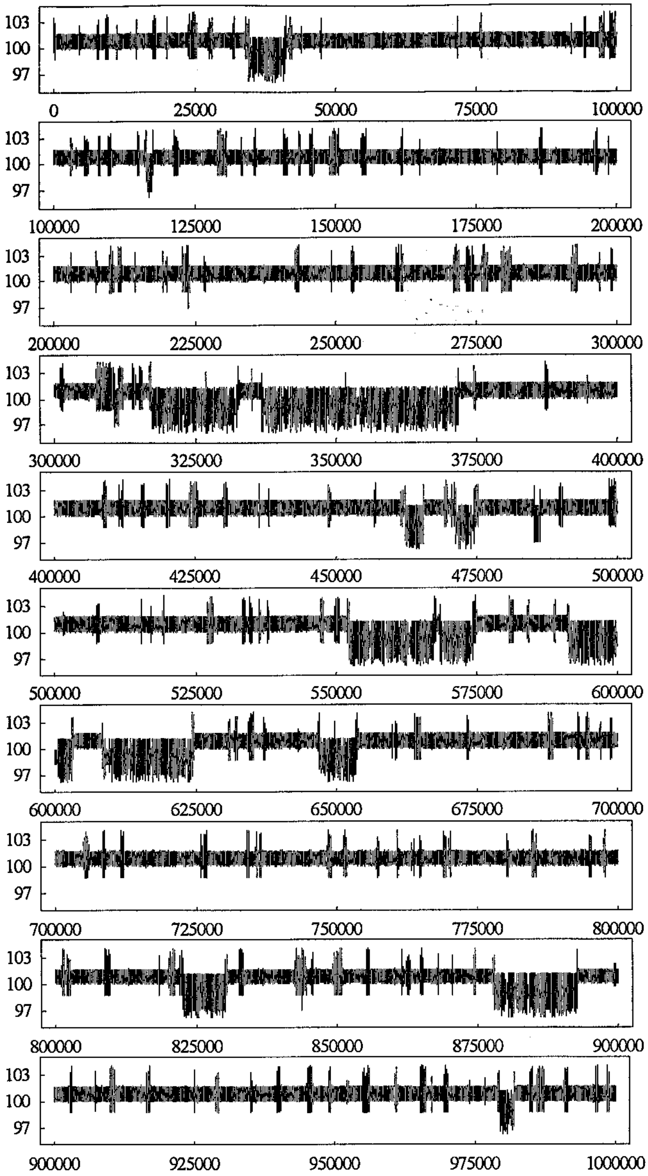

Figure 17: Changes in the Level and in the Volatility of the Exchange Rate. S^F=100, S_1=100, S_2=101, S_3=101.5, α^C=1, α^F=1, γ=0.2, β^1=0.145, β^2=170,000. 1,000,000 exchange rates are plotted in the time domain.

If the environment is extremely uncertain, the agents allow themselves to be guided by past values of the exchange rate when forming new expectations. These function as anchors in the individual judgement of the future exchange rate. This phenomenon is called anchoring heuristics and is well documented in the psychological literature (Tversky and Kahneman 1974). Then, the expectation formation is not only regressive but also anchored to the last observations of the exchange rate

$$E_t[S_{t+1}] = \gamma S^F + (1-\gamma)(S_{t-1} + S_{t-2})/2. \tag{8}$$

The demand of the fundamental trading rule modifies to

$$d_t^F = \alpha^F (E_t[S_{t+1}] - S_t)/S_t = \alpha^F (\gamma S^F + (1-\gamma)(S_{t-1} + S_{t-2})/2 - S_t)/S_t, \tag{9}$$

where the exchange rates in t-1 and t-2 are used as an orientation for the expectation formation.

The quadratic term of the weighting scheme (4) is substituted by a square root term

$$m_t = \frac{1}{1 + \beta^1 + \beta^2 \sqrt{|(E_t[S_{t+1}] - S_{t-1})/S_{t-1}|}}. \tag{10}$$

Again, the influence of the fundamentalists increases, though at a declining rate, as the relative distance between the expected future exchange rate and the spot rate rises.

Via the market clearing condition (5) the solution of the model is derived

$$S_t = \frac{\gamma S^F + (1-\gamma)(S_{t-1} + S_{t-2})/2}{1 - \dfrac{\alpha^C (0.6(LogS_{t-1} - LogS_{t-2}) + 0.4(LogS_{t-2} - LogS_{t-3}))}{\alpha^F (\beta^1 + \beta^2 \sqrt{|(\gamma S^F + (1-\gamma)(S_{t-1} + S_{t-2})/2 - S_{t-1})/S_{t-1}|})}}, \tag{11}$$

which can be simulated for specific values of the coefficients and initial values of the exchange rates.

3.4.2 Simulations

Figure 18 displays a simulation run of the exchange rate and the weight of the chartists in the time domain. Each of the modifications has a specific impact on the dynamics. Again, most of the dynamics are concentrated in a band around the fundamental value of the exchange rate. The general technical trading rule leads to a more regular up and down movement of the exchange rate. As long as the exchange rate has the tendency to move in one direction, the trading signal is not reversed. Due to the anchoring heuristic, the exchange rate remains from time to time at the turning point before reversion sets in. The square root weighting scheme cuts the number of sharp movements. These occur when the influence of chartists is strong. Due to the slope of the square root weighting scheme, the region in which the chartists strongly dominate the market is relatively small. This can also be seen in the simulated weights of chartists which vary less extremely.

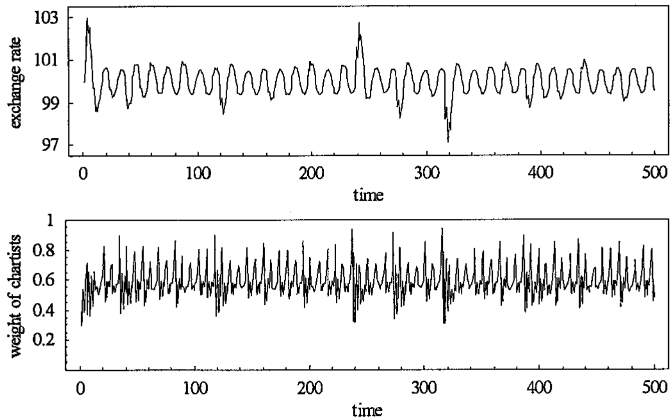

Figure 18: Exchange Rate and Weight of Chartists (GAS-Version). S^F=100, S_1=100, S_2=101, S_3=103, α^C=1, α^F=1, γ=0.2, β^1=0.05, β^2=20.

3.4.3 Quantifying the Dynamics

The bifurcation diagrams in figure 19, in which the exchange rate is plotted against β^1 (top), β^2 (middle) and γ (bottom), are similar to those of figures 4-6. An increase in β^1 or β^2 reduces the volatility. For higher β^2, the qualitative properties of the motion remain similar to those of figure 18 (in a phase space representation, the attractor simply shrinks). In the case of β^1 and γ, the dynamic behavior undergoes several qualitative changes. Fixed points, limit cycles, quasi-periodic and chaotic motion emerge.

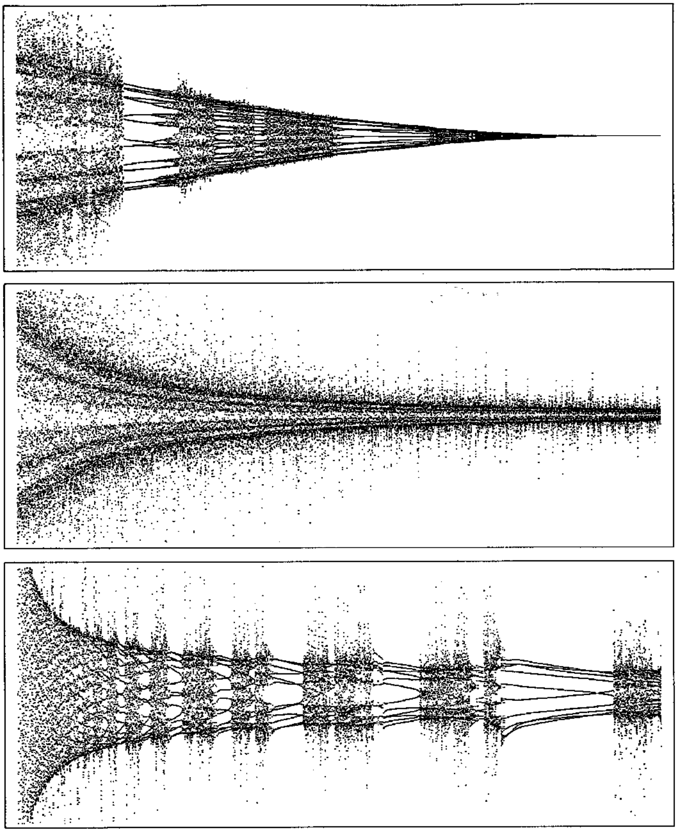

Figure 19: Bifurcation Diagrams for β^1, β^2 and γ (GAS-Version). S^F=100, S_1=100, S_2=101, S_3=103, α^C=1, α^F=1, γ=0.2, β^1=0.05, β^2=20. Top: β^1 rises in 1,000 steps from 0 to 5/7. Middle: β^2 rises from 7.5 to 32.5. Bottom: γ rises from 0.01 to 0.5. Exchange rates are plotted from t=1,001-1,050.

Figure 20 shows the evolution of the exchange rates and the returns in two-dimensional phase space. The top contains the whole attractor, the bottom demonstrates a blow-up. Although the fluctuations look more regular than in the PRQ case, again a complicated orbit structure emerges. Since the Lyapunov exponent and the correlation dimension are estimated as 0.061 and 1.627 (the base run of figure 18), we conclude that the dynamics exhibit (weak) signs of chaos.

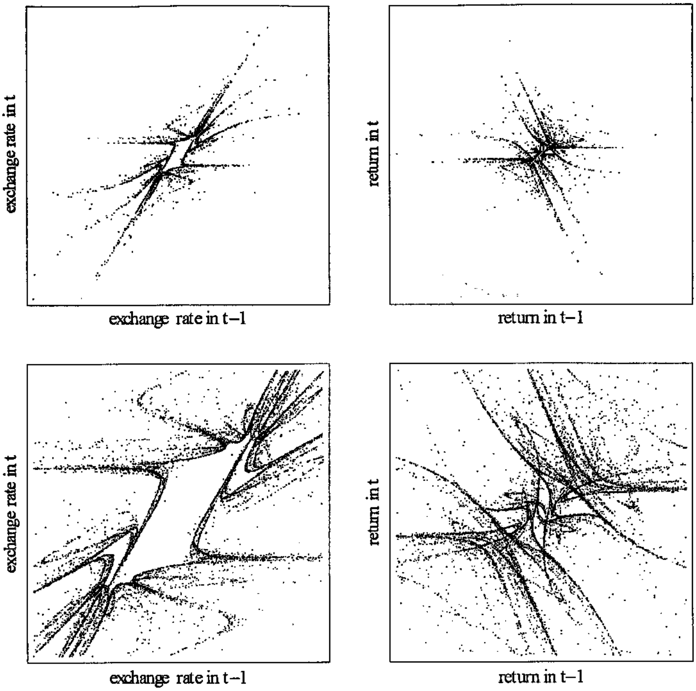

Figure 20: Exchange Rates and Returns in Phase Space (GAS-Version). S^F=100, S_1=100, S_2=101, S_3=103, α^C=1, α^F=1, γ=0.2, β^1=0.05, β^2=20. Top: data from t=1,000-11,000. Bottom: a blow-up, data from t=1,000-31,000.

Figure 21 shows that for quasi-periodic dynamics a tori emerges in phase space. In the quasi-periodic case the motion never exactly repeats itself. However, the motion is not chaotic. It is composed of two (or more) periodic components. The ratio of frequencies is irrational, that is, it cannot be expressed as a ratio of integers. Detecting the difference between quasi-periodic motion and motion with a rational ratio of frequency is a delicate question when integers are large. Moreover, it is also difficult to distinguish quasi-periodic motion from weak chaos. Each part of figure 21 contains a exchange rate time series from t=5,000-8,000 in three dimensional-phase space. The coefficient β^1 rises from the left to the right. For β^1=0.05 (the setup of figure 18), the dynamics are chaotic. However, for higher values of the coefficient, the number of outliers decreases.

For β^1=0.33, a ring (the tori) emerges. Afterwards, the behavior becomes periodic. The last part (β^1=0.36) shows a limit cycle of length 20.[8]

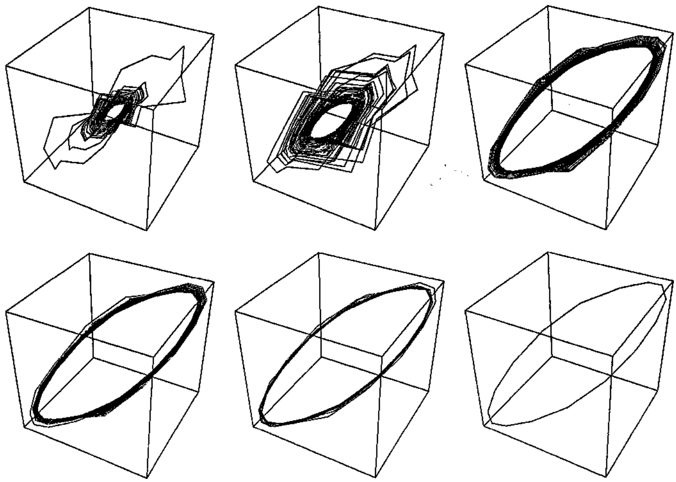

Figure 21: Phase Space Dynamics (GAS-Version). S^F=100, S_1=100, S_2=101, S_3=102.5, α^C=1, α^F=1, γ=0.2, β^2=20, β^1={0.05, 0.075, 0.26, 0.33, 0.3565, 0.36}. Exchange rates are plotted from t=5,000-8,000.

3.4.4 Robustness of the Dynamics

The model produces complex motion for different functional specifications of the underlying equations. Moreover, various parameter combinations and initial conditions lead to interesting dynamics. Next, we want to explore the stability of the model somewhat further. Remember that large exchange rate movements occur if a high demand of the chartists has to be absorbed by a low proportion of the fundamentalists. Clearly, for some parameter combinations and initial conditions one can not rule out the exchange rate becoming negative or unreasonably high. Unfortunately, due to the nonlinearity of the model we are not analytically able to derive any stability conditions.

---

[8] Of course, there are numerous qualitative changes of the dynamic behavior in between.

Nevertheless, we can investigate the stability of the model with a numerical analysis. This is done in figure 22, in which a black dot indicates a stable exchange rate path and a white dot represents an unstable one.[9] Stability is defined as follows. If the exchange rate stays for the first 50 iterations within the boundaries of $0<S<1{,}000$, we say that the motion is stable.[10]

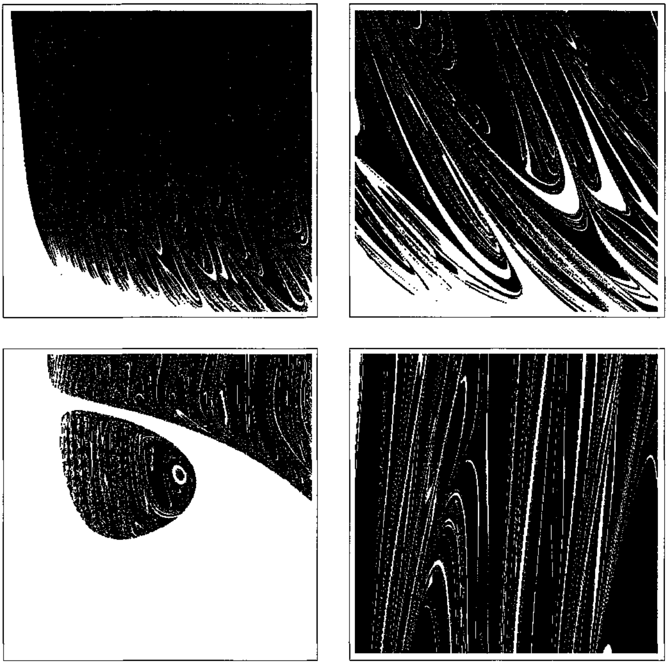

Figure 22: Basin of Attraction (GAS-Version). Top left: S^F=100, S_1=100, S_2=101, S_3=103, α^C=1, α^F=1, β^1=0.05. β^2 increases from 1.5 to 4 and γ from 0 to 0.4 (resolution 500 times 500). Top right: the same but β^2 increases from 1.5 to 2.3 and γ from 0.2 to 0.3 (resolution 400 times 500). Bottom left: S^F=100, S_1=100, S_3=103, α^C=1, α^F=1, β^1=0.05, β^2=2. S_2 increases from 1 to 100 and γ from 0 to 0.4 (resolution 400 times 400). Bottom right: the same but S_2 increases from 90 to 100 and γ from 0.2 to 0.3 (resolution 400 times 400). Black dots indicate stable orbits, defined as $0<S<1{,}000$.

[9] Since we are only interested in the stability of the trajectories, we do not test to what kind of attractor the solution streams nor how fast this would occur (as in figure 12). Clearly, not all trajectories are chaotic.

[10] Of course, both the time period and the upper bound are set at some arbitrary value. But for a high resolution, as in figure 22, computation is very time consuming. In order to speed up computation, only the first 50 iterations are calculated. For longer time periods, one would observe more unstable solutions. Nevertheless, the image of the basin of attraction remains complex.

The top section of the figure shows the basin of attraction in the β^2/γ-plane. On the right, β^2 rises from 1.5-4 and γ from 0-0.4. On the left, β^2 increases from 1.5-2.3 and γ from 0.2-0.3. The picture which emerges demonstrates why it is so difficult to determine stability boundaries analytically. Stable and unstable parameter combinations are interwoven in a very complicated way. At least the figures suggest that the dynamics are less stable if β^2 and γ decrease. But note that for low values of γ and β^2 the motion behaves unrealistically. Reasonable exchange rate fluctuations are only obtained with parameter combinations near our baseline simulation (as in figure 18, where γ=0.2 and β^2=20). In this area, we find the dynamics to be stable.

The top section of figure 22 computes stable and unstable orbits in the β^2/γ-plane. The initial conditions are always the same. The basin of attraction appears similarly intricate if one varies the initial conditions. At the bottom of the figure, the simulations are repeated in the S_2/γ-plane. At the line β^2=2, we vary S_2 on the left (right) between 1 and 100 (90 and 100). Again, a complex structure emerges.

3.5 Conclusions

The aim of this chapter is to explore the impact of speculative trading on exchange rate fluctuations within a simple, yet realistic setting. The interaction of technical and fundamental trading rules is crucial for the derived complex dynamics. In general, technical analysis tends to destabilize the market, whereas fundamental analysis has a mean reverting effect. Since the agents tend to prefer technical trading rules when the currency is not strongly misaligned, the exchange rate is driven away from its fundamental value. Due to this fact, one may draw the conclusion that financial markets are inherently instable.

Besides replicating the stylized fact of excess volatility, we have identified several sources for change in the level of the exchange rate and volatility clustering. Of course, a change in the level of the exchange rate may be caused by a permanent fundamental shock. However, in the case of coexisting attractors, a single shock may be enough for the economy to switch its dynamical behavior. Moreover, nonlinear dynamic systems even possess the ability to endogenously mimic regime shifts. As we have seen, the exchange rate may fluctuate in a certain area for some time before it suddenly shifts to another level.

Volatility clustering is partly the result of changing economic conditions. For instance, a crisis may lower the willingness of the fundamentalists to take risks. But the popularity of the trading rules may also be affected through social phenomena such as herding behavior. Transient motion of a system, triggered by some external noise, may also contribute to volatility clustering. Without assuming an arbitrary shock pattern, periods of low and high volatility may be endogenously explained by on-off intermittency.

The dynamics are robust in the sense that they are generally not affected by different functional specifications of the underlying equations, different values for the coefficients, or different initial values of the exchange rates. Of course, none of the specifications deliver a perfect fit for real-life conditions. A more reasonable interpretation would be that different specifications reflect different regimes in time. Jumping between different regimes naturally increases the complexity even further.

In this chapter we have chosen a deterministic setting to study some of the driving forces of foreign exchange dynamics. But there is no question that financial markets are affected by different kinds of stochastic disturbances. A natural extension would thus be to include some of them. To capture the variety of technical analysis one could add some random demand to these rules. In addition, the fundamental exchange rate could be assumed to behave like a random walk. Finally, although the volume of international trade transactions is relatively small, it may be an important ingredient if one aims at modeling phenomena like lasting bubbles. The incorporation of such extensions should allow us to generate time series which mimic some important stylized facts quite closely.

If this can be done in a convincing way, one can go one step further and attempt to answer some normative questions. For instance, using a chartists-fundamentalists model as a laboratory, future research may provide a greater knowledge of how central bank interventions or price controls work in practice and how they could be improved.

3.A Appendix

3.A.1 Some Notes on Nonlinear Time Series Analysis

Two key aspects of chaos are the exponential divergence of nearby trajectories and the existence of complex orbit structure. These dynamical and geometrical properties can be quantitatively characterized by the largest Lyapunov exponent and the correlation dimension, respectively. Next, we give a brief introduction to these measures. For a more complex discussion compare, for instance, Hilborn (2000), Kantz and Schreiber (1999), and Ott et al. (1994).

3.A.1.1 Lyapunov Exponent

Assume that $\mathbf{x}_n$ denotes a k-dimensional vector and that the relevant dynamical system is specified by a map[11]

$$\mathbf{x}_{n+1}=\mathbf{F}(\mathbf{x}_n)\,. \tag{12}$$

Consider an orbit displaced from the original orbit $\mathbf{x}_n \rightarrow \mathbf{x}_n+\delta\mathbf{x}_n$, where $\delta\mathbf{x}_n$ is an infinitesimal vector. The evolution of $\delta\mathbf{x}_n$ is given by

$$\delta\mathbf{x}_{n+1}=\mathbf{DF}(\mathbf{x}_n)\delta\mathbf{x}_n\,, \tag{13}$$

where $\mathbf{DF}(\mathbf{x})$ is the k×k Jacobian matrix of partial derivatives of $\mathbf{F}(\mathbf{x})$. Let $\mathbf{y}_n=\delta\mathbf{x}_n/|\delta\mathbf{x}_0|$ (the tangent vector). Then we can write

$$\mathbf{y}_{n+1}=\mathbf{DF}(\mathbf{x}_n)\mathbf{y}_n\,. \tag{14}$$

We are interested in the exponential rate at which the magnitude of $\mathbf{y}$ grows or shrinks per iteration of the map (the evolution of $\mathbf{y}_n$ depends on both $\mathbf{x}_0$ and $\mathbf{y}_0$)

$$\lambda(\mathbf{x}_0, \mathbf{y}_0, n)= (1/n)\ln|\mathbf{y}_n|\,. \tag{15}$$

Combining (14) and (15) yields

[11] For example, the PRQ version can also be written as

$$A_t = \frac{\gamma S^F + (1-\gamma)A_{t-1}}{1-(\alpha^C(0.5LogS_{t-1}+B_{t-1}))/(\alpha^F(\beta^1+\beta^2((\gamma S^F-\gamma A_{t-1})/A_{t-1})^2))},$$

$B_t = -LogA_{t-1} + C_{t-1}$,

$C_t = 0.5LogA_{t-1}$.

$$\lambda(\mathbf{x}_0, \mathbf{y}_0, n) = (1/n)\ln |\mathbf{DF}(\mathbf{x}_{n-1}) \ldots \mathbf{DF}(\mathbf{x}_0)\mathbf{y}_0|. \tag{16}$$

For $n \to \infty$, one derives the Lyapunov exponent

$$\lambda(\mathbf{x}_0, \mathbf{y}_0) = \lim_{n\to\infty} (1/n)\ln |\mathbf{DF}^n(\mathbf{x}_0)\, \mathbf{y}_0|, \tag{17}$$

where $\mathbf{F}^n$ denotes the nth iterate of $\mathbf{F}$. The Lyapunov exponent measures the exponential rate at which nearby trajectories separate. Hence, the Lyapunov number $L=e^{\lambda}$ represents an average factor by which the magnitude of the infinitesimal vector displacement is multiplied on each iterate. If λ is positive, nearby trajectories diverge exponentially. The Lyapunov exponent is negative for fixed points or periodic data, zero for quasi-periodic data, and infinite for uncorrelated random data. Besides calculating the Lyapunov exponent directly from a known system, it is also possible to estimate the Lyapunov exponent indirectly for a given time series. A prominent algorithm is described by Wolf et al. (1985).

3.A.1.2 Correlation Dimension

The correlation dimension is a measure to determine the degree of complexity of a given time series. It indicates to what degree the phase space is occupied by the attractor. For simple objects the correlation dimension is equal to their Euclidean dimension: a point has dimension 0, a line has dimension 1, and a plane has dimension 2. More complicated objects (say a cloud) often have noninteger dimensions. The figure in phase space is then called a strange attractor. The correlation dimension is derived from the correlation sum. The correlation sum for a collection of points $\mathbf{x}_n$ in some vector space is defined as the fraction of all possible pairs of points which are closer than a given distance ε in a particular norm

$$C(\varepsilon) = \frac{2}{N(N-1)} \sum_{i}^{N} \sum_{j=i+1}^{N} \Theta(\varepsilon - \| \mathbf{x}_i - \mathbf{x}_j \|), \tag{18}$$

where $\Theta(x)=0$ for $x \leq 0$ and $\Theta(x)=1$ for $x>0$. The sum counts the pairs of points $(\mathbf{x}_i, \mathbf{x}_j)$ whose distance is smaller than ε. In the limit of an infinite amount of data and for small ε, the correlation dimension C scales according to a power law $C(\varepsilon) \approx \varepsilon^D$. Hence, the correlation dimension is obtained as

$$D = \lim_{N\to\infty}\lim_{\varepsilon\to\infty} \frac{\partial \ln C(N,\varepsilon)}{\partial \ln \varepsilon}. \tag{19}$$

The dimensionality of an attractor indicates the active degrees of freedom for the system, that is, the number of variables needed to obtain the object. An upper bound of the correlation dimension is given by the dimension of the system. If experimental data shows a low dimensional attractor (D<4), one may be able to understand the underlying dynamics. However, a higher dimension estimate indicates that the dynamics are close to random.

3.A.2 Summary of all Versions

Figures 23 – 30 present a graphical summary of the dynamics and contain the exchange rate in time domain and in phase space, the return in phase space, and a bifurcation diagram. Table 1 contains estimates of the Lyapunov exponent and the correlation dimension. The underlying parameter settings correspond to those used for the simulation of the exchange rate in time domain. The software employed for data analysis is Chaos Data Analyzer: The Professional Version 2.1 by Sprott and Rowlands (1995a). Information regarding the description of statistics are obtainable from their PC user´s manual (Sprott and Rowlands 1995b). The algorithm to measure the Lyapunov exponent is due to Wolf et al. (1985) and the algorithm to estimate the correlation dimension is a variation of that by Grassberger and Procaccia (1984).

To test whether the evidence of hidden determinism in the data sets is robust, it is prudent to repeat the calculations of the tests using surrogate data that resemble the original data but with the determinism removed. Robustness implies that analysis of the surrogate data should provide values that are statistically distinct from those calculated from the original data. One way to check the results is simply to shuffle the data values. This operation preserves the probability distribution but generally produces a very different power spectrum and correlation function. A second approach is to apply a Fourier transformation, randomize the phase of each Fourier component, and then invert the Fourier transformation. This operations preserves the power spectrum and correlation function but generally produces a different probability distribution. We find that for both tests the estimated values lie outside the range of values calculated from the original data sets and are thus really chaotic.

Summary: From figures 23-30 and table 1 we draw the conclusion that all eight versions have the potential to produce complex dynamics. Thus, the dynamics are generally robust for different functional specifications, or values of the initial conditions and parameters. The PRS, GRQ, PRQ, GAQ and PAQ versions exhibit clear signs of chaos. The GRS, GAS and PAS versions only show small estimates of the Lyapunov exponents (the time series behavior is somehow on the edge between quasi-periodic and weak chaotic motion).

version	GRS	PRS	GAS	PAS
Lyapunov exponent	0.009	0.154	0.061	0.008
correlation dimension	0.921	2.131	1.627	1.014
version	GRQ	PRQ	GAQ	PAQ
Lyapunov exponent	0.152	0.142	0.189	0.265
correlation dimension	1.186	1.620	1.655	1.819

Table 1: Chaos Data Analysis. For the Lyapunov exponent the calculations considered the proper embedding dimension as given by 3, using 3 time steps and an accuracy of 10^{-4}. For the correlation dimension the calculations considered the proper embedding dimension equal to 5 and time delay equal to 1. Each calculation is based on 10,000 data points. The parameter values underlying the simulation correspond to the specifications used in Appendix 1a - 2h. The versions are coded as follows: G=general and P=popular technical trading rule, R=regressive and A=anchor expectations, S=square root and Q=quadratic weighting scheme.

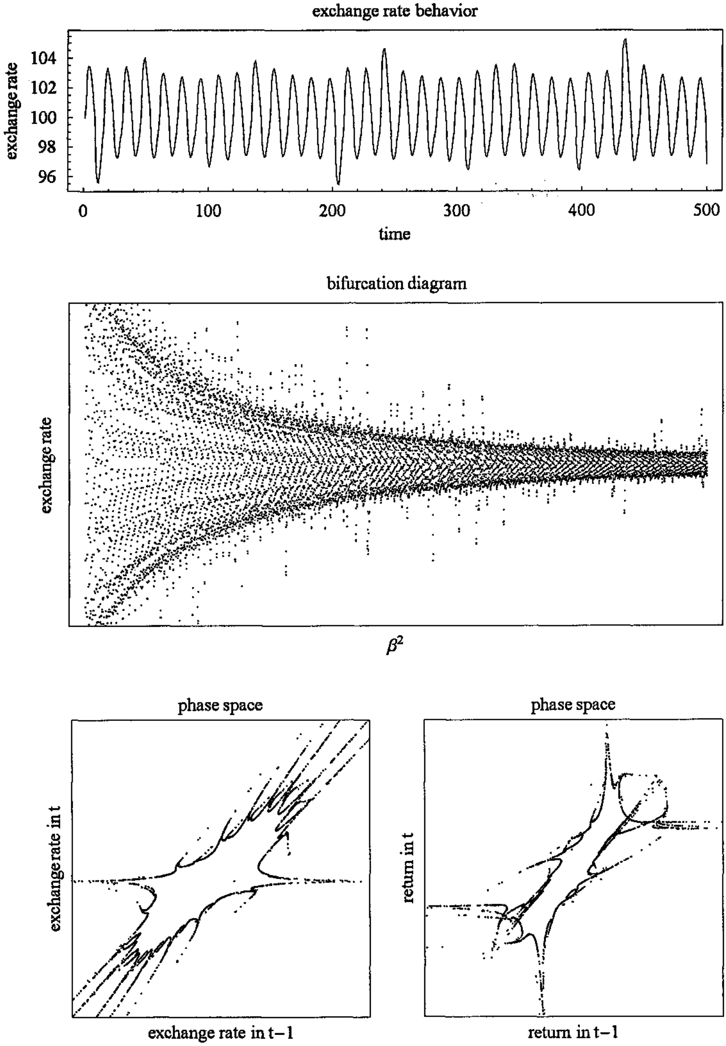

Figure 23: GRS-Version. $S^F=100$, $S_1=100$, $S_2=101$, $S_3=103$, $\alpha^C=\alpha^F=1$, $\gamma=0.2$, $\beta^1=0.05$, $\beta^2=20$. Bifurcation diagram: β^2 rises from 15 to 65.

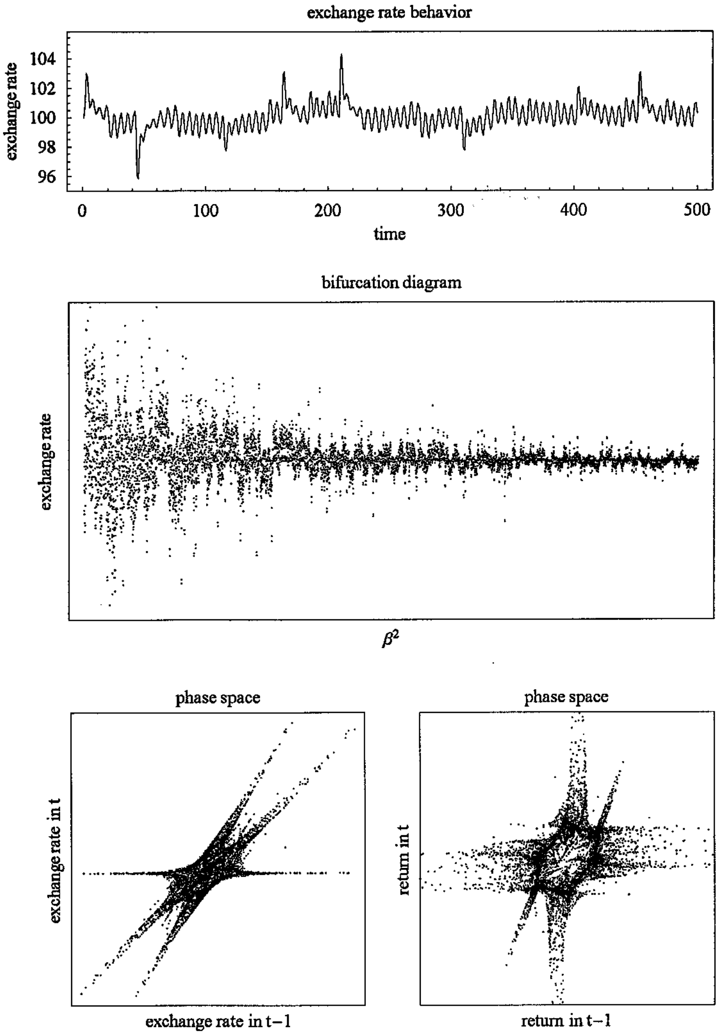

Figure 24: PRS-Version. S^F=100, S_1=100, S_2=101, S_3=103, α^C=α^F=1, γ=0.2, β^1=0.075, β^2=40. Bifurcation diagram: β^2 rises from 30 to 130.

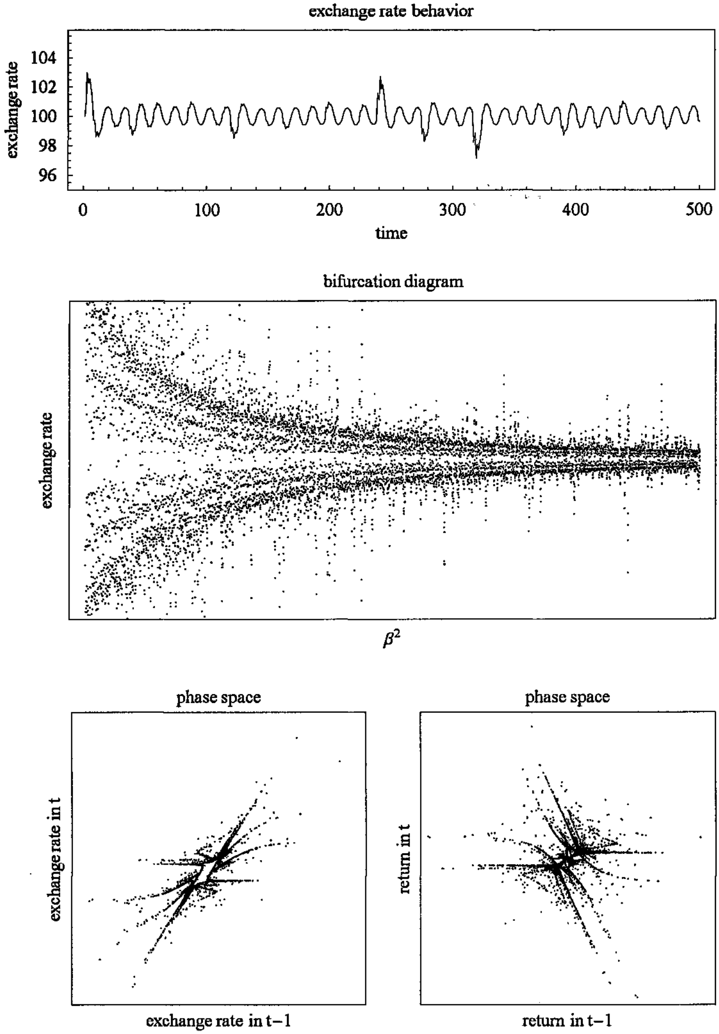

Figure 25: GAS-Version. S^F=100, S_1=100, S_2=101, S_3=103, α^C=α^F=1, γ=0.2, β^1=0.05, β^2=20. Bifurcation diagram: β^2 rises from 7.5 to 32.5.

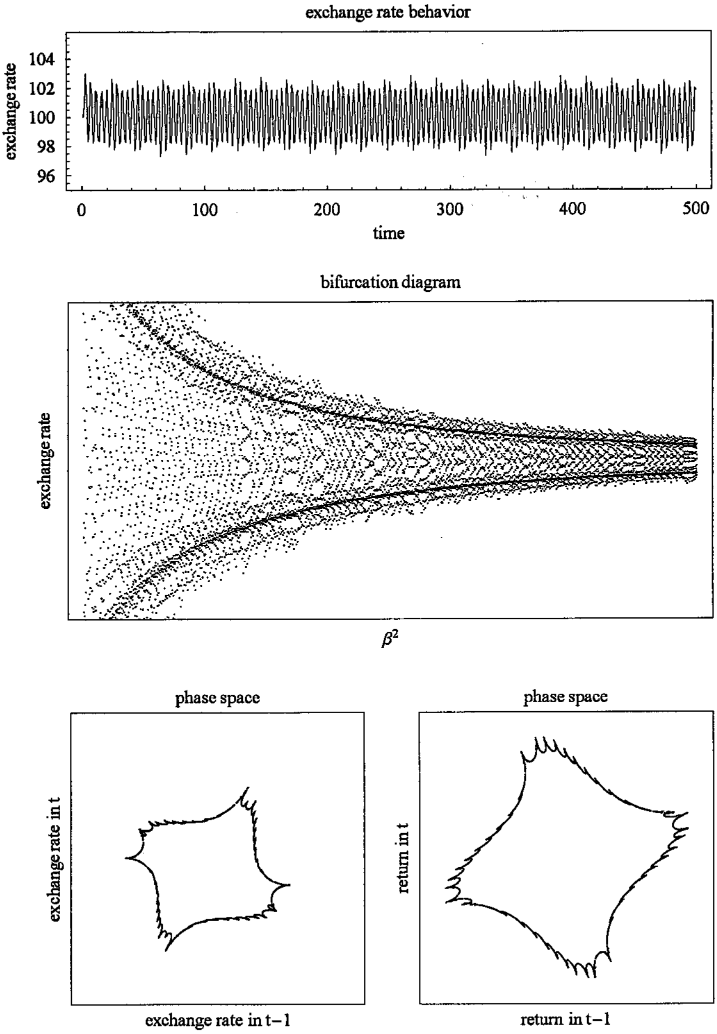

Figure 26: PAS-Version. $S^F=100$, $S_1=100$, $S_2=101$, $S_3=103$, $\alpha^C=\alpha^F=1$, $\gamma=0.6$, $\beta^1=1$, $\beta^2=6$.
Bifurcation diagram: β^2 rises from 3.5 to 13.5.

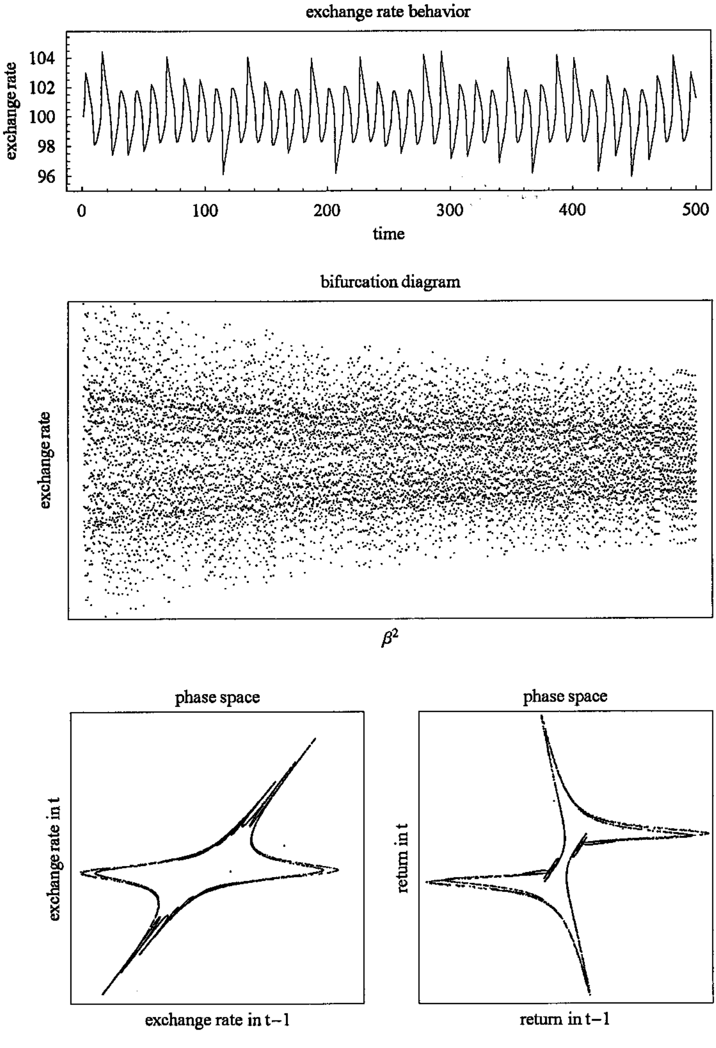

Figure 27: GRQ-Version. $S^F=S_1=100$, $S_2=101$, $S_3=103$, $\alpha^C=\alpha^F=1$, $\gamma=0.2$, $\beta^1=0.15$, $\beta^2=25{,}000$. Bifurcation diagram: β^2 rises from 300,000 to 1,300,000.

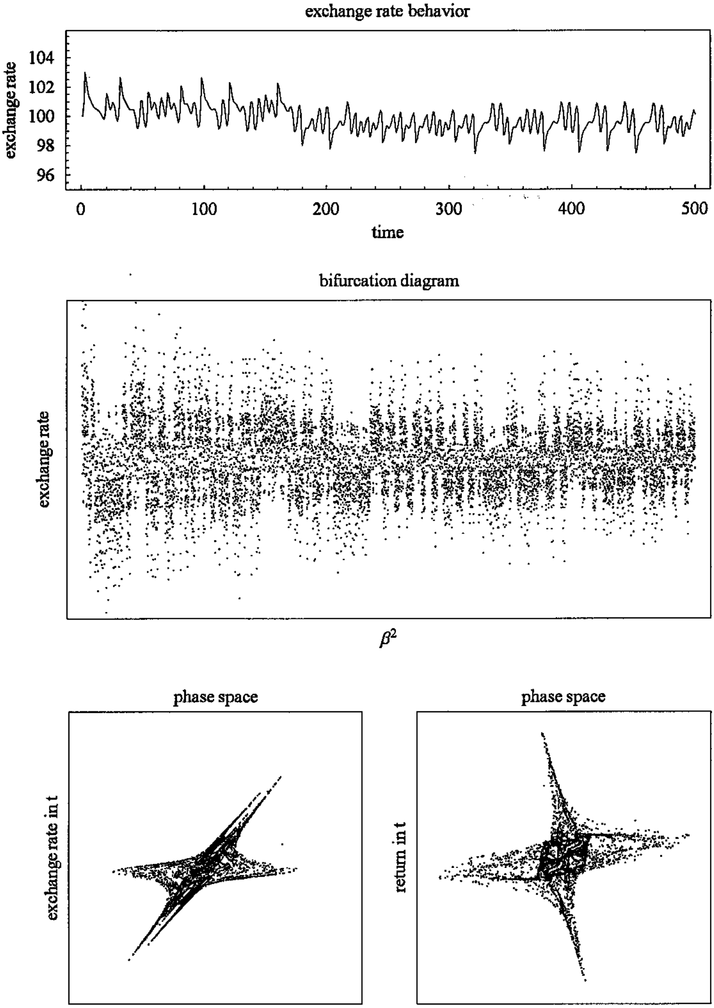

Figure 28: PRQ-Version. $S^F=S_1=100$, $S_2=101$, $S_3=103$, $\alpha^C=\alpha^F=1$, $\gamma=0.2$, $\beta^1=0.25$, $\beta^2=640{,}000$. Bifurcation diagram: β^2 rises from 300,000 to 1,300,000.

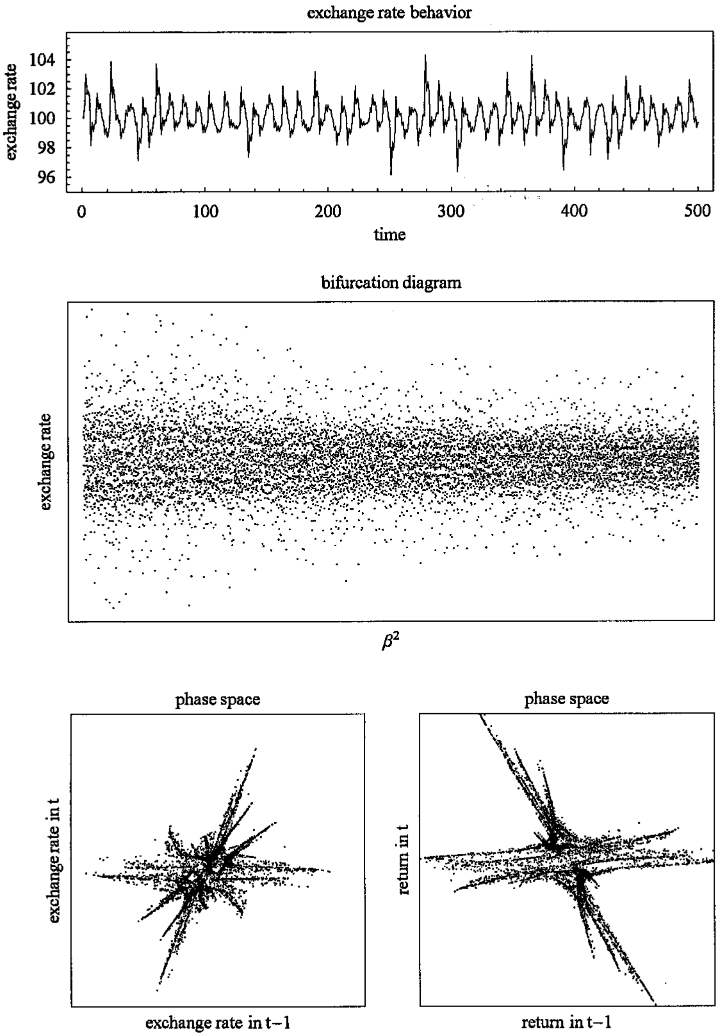

Figure 29: GAQ-Version. $S^F=S_1=100$, $S_2=101$, $S_3=103$, $\alpha^C=\alpha^F=1$, $\gamma=0.2$, $\beta^1=0.075$, $\beta^2=160{,}000$. Bifurcation diagram: β^2 rises from 150,000 to 650,000.

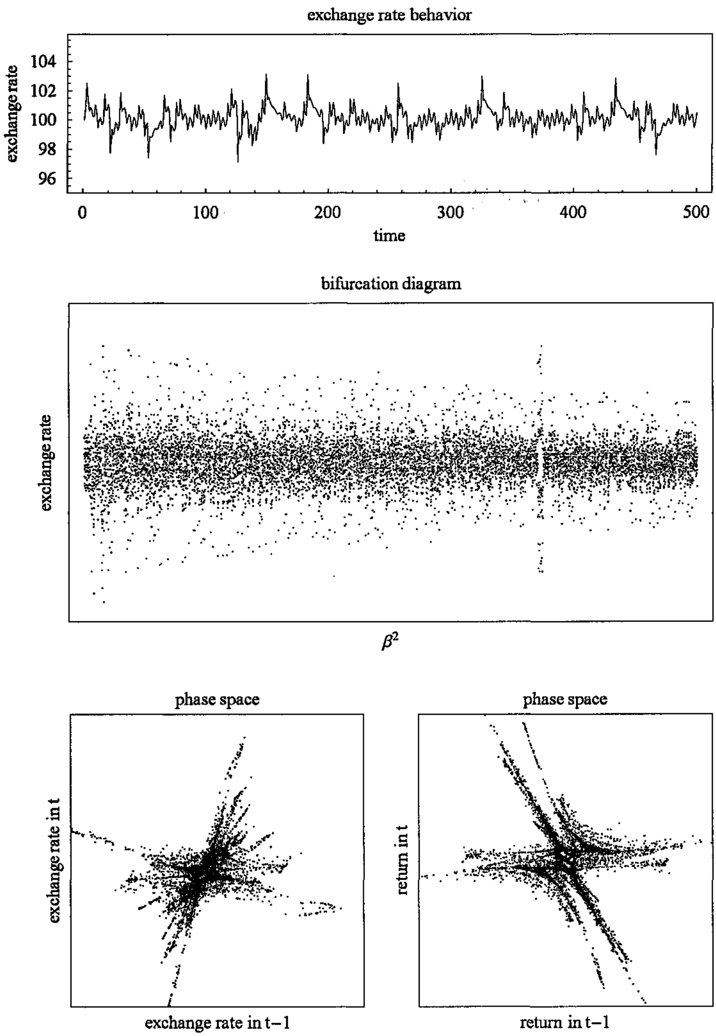

Figure 30: PAQ-Version. $S^F=S_1=100$, $S_2=101$, $S_3=103$, $\alpha^C=\alpha^F=1$, $\gamma=0.2$, $\beta^1=0.25$, $\beta^2=1{,}000{,}000$. Bifurcation diagram: β^2 rises from 300,000 to 3,900,000.

4 Speculative Behavior, Exchange Rate Volatility, and Central Bank Intervention

Abstract

The aim of this chapter* is twofold. First, to develop a model which helps to explain the high exchange rate volatility observed empirically. Second, to study under which conditions central bank interventions may calm down the foreign exchange market. Based on empirical observations, a model is presented where the agents select in each trading period a trading rule to determine their speculative positions. The agents have the choice between technical and fundamental trading rules. Simulations produce a high variability of the exchange rates, fat tails for returns, and weak evidence of mean reversion. Within this framework, the effectiveness of some intervention strategies is analyzed. One result is: "leaning against the wind" may reduce the volatility as long as the dynamics are influenced by trend-following trading strategies. In periods when the agents are uncertain about the fundamental exchange rate, however, supporting a target exchange rate may be the preferable strategy for the central bank to stabilize the market.

Keywords

exchange rate theory, technical and fundamental trading rules, central bank intervention

JEL Classification

E58, F31, G14

* This chapter was presented at the ZEI-Summer School in International Finance 2000, Bad Honnef, August 2000 and at the 7th Viennese Workshop on Optimal Control, Dynamic Games and Nonlinear Dynamics, Vienna, May 2000. I would like to thank the participants for their helpful discussions, especially Carl Chiarella, Giancarlo Corsetti, Jürgen von Hagen and two anonymous referees.

4.1 Introduction

Since the development of real time information systems and the decline in transaction costs following the liberalization of the capital markets in the mid 80's, both daily foreign exchange turnover and volatility of exchange rates have sharply increased. The trading volume reflects more and more short-term transactions indicating a highly speculative component. By contrast, international trade transactions account for merely one percent of the total (BIS 1999).

When determining their speculative positions, the market participants rely on both technical and fundamental trading rules. Technical analysis is a trading method that attempts to identify trends and reversals of trends by inferring future price movements from those of the recent past whereas fundamental concepts look at the underlying reasons behind that action.[1] As reported by Taylor and Allen (1992) most foreign exchange dealers place at least some weight on technical analysis, especially in the short run.

Based on such evidence, the noise trader approach (Shleifer and Summers 1990) is a research direction which focuses on modeling speculative behavior. The noise trader approach assumes that not all market participants are fully rational and that arbitrage possibilities are limited. Consequently, shifts in investor sentiment are an important determinant of high exchange rate volatility. The more specific chartists-fundamentalists models are a special branch within this research program (among them Frankel and Froot 1986, de Grauwe et al. 1993, and Lux 1997). Central for these papers is the interaction between chartists (technical traders) and fundamentalists. This area of research is very promising because some basic stylized facts of the empirical data are replicated.

The aim of this chapter is twofold. First, to develop a realistic, yet simple exchange rate model in the spirit of the chartists-fundamentalists models to get a deeper understanding of the driving forces behind foreign exchange dynamics. Second, to evaluate whether

[1] For an introduction into technical analysis see Neely (1997). A deeper discussion is found in Murphy (1999). The latter is sometimes referred to as the "bible" of technical analysis.

typical intervention strategies of central banks will be able to calm down disorderly markets dominated by speculative transactions.

The model presented in this chapter is similar to models found in the chartists-fundamentalists literature. Rather than deriving the results from a well defined utility maximization problem, details from the market microstructure and psychological evidence are used to describe the behavior of the traders. Clearly, the traders rely on simple decision rules to determine their investment positions. The interaction of these rules generates, even in a simple setting, a realistic behavior of exchange rates. In their first moments the simulated time series resemble a stochastic trend. Evidence of mean reversion is observable if the behavior of the chartists is trend-following. The returns of the generated time series display a high kurtosis which declines under time aggregation. Fat tails are also identified by the scaling behavior of the returns which roughly follow a power law.

The volatility is mostly caused endogenously through the interaction between the traders. The news arrival process plays a role in the sense that changes in fundamentals are amplified by technical trading. Although the dynamics in the foreign exchange market are very complex, simple nonlinear models may suffice to explain and understand them. Such models preclude predictions but invite to investigate how the underlying system may be controlled. Thus, we introduce a central bank into the model and analyze the effectiveness of some well-known intervention strategies. Note that we are not searching for an optimal intervention strategy. Rather we ask what the conditions for a successful intervention operation are and when it is likely that the intervention operation fails. The results are: if the dynamics are dominated by trend-following trading strategies, a "leaning against the wind" intervention where the central bank operates against the trend is able to reduce the volatility. The strategy becomes less successful when the chartists trade rather unsystematically. However, in periods where the fundamentalists are uncertain about the fundamental value of the exchange rate "leaning against the wind" is likely to increase the variability of the market. This explains why in empirical studies the effectiveness of central bank intervention changes in sign across time. In such a situation the central bank may be better off by directly supporting the target rate.

This chapter is organized as follows. Section 4.2 presents the model and derives some simulation results. In section 4.3, a central bank is introduced into the model to study the consequences of some intervention strategies. Section 4.4 offers some conclusions.

4.2 A Simple Nonlinear Exchange Rate Model

4.2.1 Setup of the Model

The basic idea of the model is that the market participants have to choose at the beginning of each trading period a specific trading rule to determine their speculative positions. The selection of the rules depends somehow on expected performance possibilities, which are derived from past observations. The agents have the choice between technical and fundamental trading rules. The former are called chartists and the latter fundamentalists.

Simple technical trading rules use only past movements of the exchange rate S as an indicator of market sentiment and extrapolate these into the future, thus adding a positive feedback to the dynamics. The excess demand of chartists in period t resulting from such rules might be expressed as

$$d_t^C = a^{C,1}(0.6(LogS_{t-1} - LogS_{t-2}) + 0.4(LogS_{t-2} - LogS_{t-3})) + a^{C,2}\delta_{t-1}. \qquad (1)$$

The first bracket of (1) describes a simple moving average rule to capture the usual behavior of the chartists. In general, chartists buy (sell) foreign currency if the exchange rate rises (declines). Since more attention is paid to the most recent trend, a larger coefficient is selected for the first extrapolating term than for the second term (0.6 versus 0.4). The second bracket represents additional random demand to allow for more complicated behavior, where δ is an independently and identically distributed normal random variable with mean zero and time invariant variance. With the (positive) coefficients $a^{C,1}$ and $a^{C,2}$ the relation between the systematic and unsystematic demand components is calibrated. For simplicity, these coefficients are not time dependent. Note that by (1) chartists place a market order today in response to past price changes, i.e. 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).

Fundamental trading rules depend on the expected future exchange rate. The expectation formation process of the agents is modeled in a typically regressive way, i.e. when the exchange rate deviates from its equilibrium value S^F, the fundamentalists expect it to return. Therefore, $E_t[S_{t+1}]=\gamma S^F_{t-1}+(1-\gamma)S_{t-1}$, where γ stands for 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 excess demand of fundamentalists can be written as

$$d_t^F = \alpha^F (E_t[S_{t+1}] - S_t)/S_t = \alpha^F (\gamma S_{t-1}^F + (1-\gamma)S_{t-1} - S_t)/S_t, \tag{2}$$

where α^F is a positive reaction coefficient. The fundamental trading rule delivers a buy (sell) signal, if the expected future exchange rate is above (below) the spot rate. The corresponding demand depends on the relative distance between the expected rate and the spot rate.[2]

The agents form their expectations of the fundamental exchange rate on the basis of a structural model. The development 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

$$LogS_t^F = LogS_{t-1}^F + p\varepsilon_t. \tag{3}$$

The news ε_t (the jump size) is identically and independently distributed according to a Normal distribution with mean zero and time invariant variance. The news hits the market with $prob\ (p=1) = 0.2$ (the jump arrival time intensity). On average, a shock hits the market every 5 periods.

The selection of the rules depends on expected future performance possibilities. Fundamentalism, compared to chartism, becomes more popular the wider the spot rate deviates from the expected future exchange rate. This might be justified as follows. The chance that the exchange rate returns to its expected value increases as its relative distance rises. We define the weight of chartists as

[2] Due to the time structure of the model the fundamentalists function as market makers.

$$m_t = 1/(1+\beta^1 + \beta^2 \sqrt{|(E_t[S_{t+1}] - S_{t-1})/S_{t-1}|}, \tag{4}$$

and the weight of fundamentalists as $(1-m_t)$, respectively. The coefficient β^1 represents the basic influence of the fundamentalists. If, for example, β^1 is 0.25, then 20 percent of the agents are always fundamentalists. Nevertheless, most traders adjust their trading strategies with respect to the relevant conditions. As assumed by (4), the weight of fundamentalists increases, though at a declining rate, as the relative distance between $E_t[S_{t+1}]$ and S_{t-1} rises. In such a situation, more and more agents realize that the exchange rate does not reflect its fundamental value any more. Consequently, fundamental analysis is preferred to chartism. Note that the influence is determined with a time lag since the selection of the rules has to be repeated at the beginning of each new trading period.

Demand from international trade is neglected since trade transactions, in contrast to speculative transactions, are small in absolute magnitude. Using the market clearing condition

$$m_t d_t^C + (1-m_t) d_t^F = 0, \tag{5}$$

the solution of the model is a four-dimensional stochastic difference equation system

$$S_t = \frac{\gamma S_{t-1}^F + (1-\gamma) S_{t-1}}{1 - \dfrac{a^{C,1}(0.6LogS_{t-1} - 0.2LogS_{t-2} - 0.4LogS_{t-3}) + a^{C,2}\delta_{t-1}}{a^F(\beta^1 + \beta^2 \sqrt{|(\gamma S_{t-1}^F + (1-\gamma)S_{t-1} - S_{t-1})/S_{t-1}|})}}. \tag{6}$$

Since (6) cannot be solved explicitly, some simulations are done to demonstrate that the underlying structure gives rise to complex exchange rate behavior as it is typically observed empirically.[3, 4]

[3] If a low proportion of fundamentalists is confronted with a huge demand of chartists, a large price reaction is needed in order to match the demand. However, a stability problem in (6) never occurred in our simulations.

[4] Chapter 3 shows that in the absence of any shocks, i.e. $\delta=\varepsilon=0$, the model can generate chaotic motion (positive Lyapunov exponents, low dimensional attractors). This finding is observed for different parameter settings and functional specifications of (1), (2), and (4). The mechanics of the system are best described by an endogenous stretching and folding of the exchange rate around its fundamental value. Although the dynamics are very complex, we allow for some shocks in order to mimic empirical exchange rate fluctuations more closely.

4.2.2 Simulations

The top of figure 1 displays the simulated dynamics for the exchange rate (solid line) and the stochastic development of its fundamental (dashed line), the bottom presents the weights of chartists. Even a low probability of fundamental shocks suffices to generate complex exchange rate movements, where the exchange rate fluctuates around its fundamental value. Moreover, the volatility of the exchange rate is far greater than the volatility of its equilibrium value. The influence of chartists is concentrated in the range from 40 to 60 percent with some peaks going down to 25 or up to 80 percent. Such a behavior is pretty close to what is reported in survey studies (Taylor and Allen 1992).

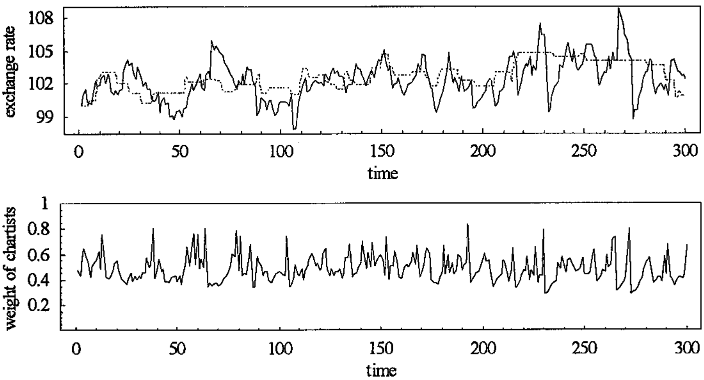

Figure 1: Simulated Exchange Rates and Weights of Chartists. The solid line in the top is the exchange rate, the dashed line its fundamental, $S_1^F=100$, $S_1=100$, $S_2=101$, $S_3=101.5$, $a^{C,1}=0.2$, $a^{C,2}=0.0064$, $a^F=1$, $\beta^1=0.1$, $\beta^2=22$, $\gamma=0.2$, $\delta \sim N(0,1)$, $prob(p=1)=0.2$, $\varepsilon \sim N(0, 0.0075)$, T=300.

Simplified, the dynamics can be explained as follows. Technical trading rules always produce some kind of buy or sell signal. On the basis of a feedback process, a self-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 stronger the influence of the fundamentalists, until eventually their increasing net position triggers a mean reversion. However, this indicates a new signal for the chartists and directly leads to the next momentum. The exchange rate overshoots the fundamental exchange rate because chartism dominates the market near the fundamental. Heavy outliers occur when the chartists have a clear trading signal and the influence of the fundamentalists is low. Since the exchange rates move several periods in one direction, chartism may be profitable temporarily.

Note that the simulations indicate that the volatility of the foreign exchange market need not be solely caused by exogenous shocks; it might be explained at least partially by an endogenous nonlinear law of motion. The trading signals needed to keep the process going are generated by the agents themselves.

A lot of empirical work is done on describing the distribution of the returns. Figure 2 compares the distribution of the returns and the scaling behavior for the simulated data (top) with normally distributed returns (bottom). An important stylized fact says that the distribution of the returns reveals fat tails (Guillaume et al. 1997). In contrast 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 are able to reveal fat tails. Table 1 displays estimates of the kurtosis under time aggregation for 20,000 data points. In comparison, the kurtosis of a Normal distribution is given with 3. Since the random variables are normally distributed, the high kurtosis is caused through the model. Stronger outliers do not only occur as a consequence of normally distributed shocks. If, for instance, a medium demand by chartists is matched by a low weight of fundamentalists, the price reaction is also strong. Furthermore, the empirically observed kurtosis declines under time aggregation.[5] This is also true for the kurtosis of the computed time series.

time aggregation	1	2	5	10	25	50
kurtosis	12.7	6.3	3.9	3.2	3.4	3.3

Table 1: Kurtosis under Time Aggregation. The same parameter setting as in figure 1, T=20,000.

An alternative way to identify fat tails is to determine the tail index. The tail index α, given as $F(|return|>x) \approx cx^{-\alpha}$, is estimated from the cumulative distribution of the positive and negative tails for normalized Log-returns. The returns are normalized by dividing

[5] A time aggregation of *d* means that the returns are calculated as $r_t = LogS_t - LogS_{t-d}$.

by the standard deviation. A regression on the largest 30 percent of the observations delivers a significant tail index of 3.3 which is in good agreement with results obtained from empirical data. According to Lux and Ausloos (2000), the tail index has mostly been found to hover between 2.5 and 4. The tail index of a Normal distribution, as can be seen at the slope in the bottom right part of figure 2, is clearly higher.[6]

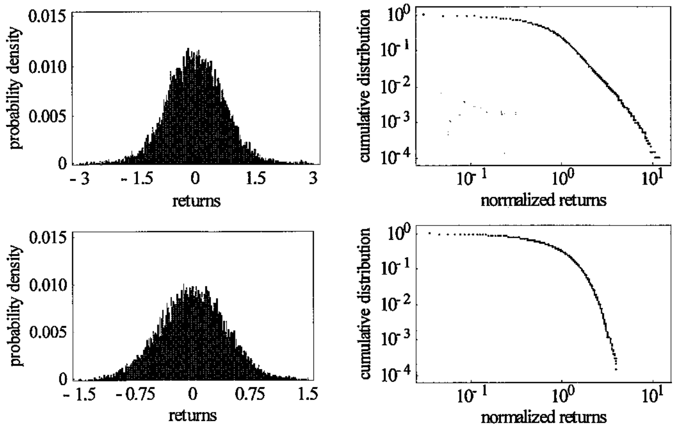

Figure 2: Distribution of Returns and Scaling Behavior. The top contains the distribution of returns and the scaling behavior of the cumulative distribution of the positive and negative tails for normalized Log-returns, the bottom the same for a Normal distribution with identical variance, the same parameter setting as in figure 1, T=20,000.

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. Figure 3 displays the autocorrelation function of the returns for three different numerical specifications of equation (1): in the top the systematic demand is roughly 50 percent of total transactions, in the middle 30 percent, and in the bottom 20 percent. The bottom part contains the autocorrelation function for the earlier simulations. Depending on the extent to which the demand of the chartists is correlated, the simulated time series reveals some kind of mean reversion tendency. Clearly, the

[6] Note that the moments of a distribution higher than its tail index are not bounded (Guillaume et al. 1997).

empirically observed autocorrelation function lies somehow between the ones shown in the middle and the bottom part of the figure. 95 percent confidence intervals are given as $\pm 2/\sqrt{T}$, with T as the number of observations and the assumption of white noise of the returns.

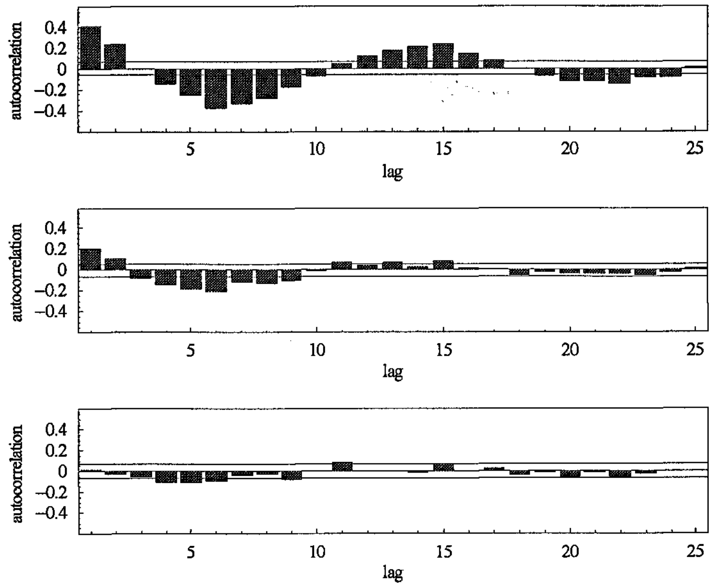

Figure 3: Autocorrelation Function of Returns. The same parameter setting as in figure 1, but in the top $a^{C,1}$=0.75, $a^{C,2}$=0.0045, in the middle $a^{C,1}$=0.45, $a^{C,2}$=0.0055, and in the bottom $a^{C,1}$=0.2, $a^{C,2}$=0.0064, T=1,000, 95 percent confidence intervals are plotted as $\pm 2/\sqrt{T}$ (assumption of white noise).

All in all, the model presented generates complex exchange rate movements and replicates some well-known stylized facts. On the one hand, the computed time series looks apparently random. On the other hand, some (deterministic) pattern like mean reversion is also observable. These features are the outcome of the nonlinear structure in the model. Evidence of nonlinearities in financial data is, for instance, strongly supported by Barnett and Serletis (2000). We argue that the forces driving the dynamics are at least partially endogenous. Although the dynamics are highly complex, simple models may suffice to explain and understand them. Such models preclude predictions but invite to control the underlying system.

4.3 The Effectiveness of Central Bank Intervention

4.3.1 Central Bank Intervention

Central bank intervention is the practice of monetary authorities buying and selling currency in the foreign exchange market to influence the exchange rate. Our focus will be on sterilized interventions and not on ordinary monetary policy. Sterilized interventions are intervention operations that are accompanied by an offsetting open market operation that restores the domestic monetary base to its original size.[7]

One aim of central bank interventions could be to influence the level of the exchange rate, for instance to achieve policy goals or to limit misalignments of the exchange rate. There are two channels through which official interventions might affect the foreign exchange market. Under the assumption that domestic and foreign assets are imperfect substitutes, the portfolio balance approach states that investors allocate their portfolios to balance their risk against expected rates of returns. If intervention operations change the relative supplies of assets denominated in different currencies, investors must be compensated to hold the relatively more numerous asset with a higher expected return. This higher expected return must result from a change in either the price of the asset or the exchange rate (portfolio balance channel). Through the signaling channel, sterilized interventions can have an effect on exchange rates if the interventions provide the market with relevant information previously not known or not fully incorporated in the current exchange rate. However, there is little empirical evidence that interventions have an influence on the level of the exchange rate.

Another aim of central bank interventions could be to reduce the exchange rate volatility, since exchange rate risk has a negative impact on the international trade. Although financial markets provide some hedge instruments, these only allow a limited elimination of the risk. A necessary condition for volatility decreasing intervention is that the variability is caused at least partially endogenously and is not solely justified by exogenous shocks. Then, intervention might work through the noise trader channel. In

[7] For surveys on central bank interventions see Edison (1993) or Dominguez and Frankel (1993).

our model the exchange rate is determined by the demand and supply flowing through the foreign exchange market. These transactions are considerably affected by noise trader activity. In such a short-run flow equilibrium, a central bank is able to manipulate the exchange rate at least at the moment the intervention takes place. Since the chartists assign much heavier weight to the most recent exchange rate movements when taking positions, the effect of the intervention is not only transitory by itself, but also amplified and prolonged by noise trader activity. Through this channel, a central bank intervention can slow down the momentum of an exchange rate trend or even reverse the direction of the trend.[8]

The noise trader channel may also explain the high degree of secret interventions. As mentioned by Murphy (1999), the philosophy behind technical trading rules is that if prices rise, the fundamentals must be "bullish". Thus, technicians claim to study the fundamentals indirectly. As long as it is unknown that the central bank is responsible for an exchange rate movement, chartists believe in the trend. If they know that the market reaction is caused by an intervention, they will become suspicious and may even counter it.

Until recently, empirical research examining the impact of sterilized intervention on exchange rate volatility barely existed. However, a promising study from Hung (1997) identifies periods where intervention operations significantly decreased the volatility via the noise trader channel. In addition, periods where intervention operations have increased the volatility are also detected. According to Hung, these operations have not to be unsuccessful intervention periods, since the central banks may have used, in order to reach a target exchange rate, the chartists by inducing a momentum in the desired direction. But volatility-enhancing operations, compared to volatility-decreasing operations, are less often used. Nevertheless, the noise trader channel allows the explanation of the phenomenon that the impact of intervention operations on the exchange rate volatility changes signs across time.

[8] To evaluate the welfare implications correctly, one has to take into account the costs of intervention operations. But as empirical studies show, US intervention operations were cost free in the past (Neely 1997, LeBaron 1999).

In the following we try to develop a deeper understanding of some intervention mechanism in the noise trader framework by introducing a central bank into the model. To start, we look at how intervention operations are executed in practice. Recently some studies became available which have access to daily intervention data (Neely 1998, LeBaron 1999, and Saacke 1999). The decision of central banks whether and how to intervene seems to be made on a day to day basis. For example, both the Federal Reserve Bank and the Bundesbank intervened in the period between 1979 and 1996 on one day in four. The interventions tend to be clustered together in time. The probability of intervention for a day strongly increases if there has been an intervention the day before. If intervention did occur, it was small – in absolute value – relative to the size of total transactions. Nevertheless, at the very moment the intervention takes place its volume hits a considerably thinner market. In addition, the interventions are typically sterilized, on average relatively balanced and performed secretly.

Finally, regression studies about the intervention reaction function of a monetary authority indicate that interventions are significantly influenced by past changes in the spot exchange rate or by deviations of the exchange rate from a target rate. More clearly, the central banks engage in so called “leaning against the wind” operations, that is, they buy (sell) foreign currency if the exchange rate declines (rises) in order to reduce the momentum of a trend, or the central banks intervene to support a target exchange rate.

4.3.2 Strategies and Goals

In this section the intervention strategies and the aims of the central bank are formalized. We concentrate on the two most common strategies empirically identified. Furthermore, we assume that interventions take place every period and that the decision about the intervention volume for period t has to be made before the trading starts. This seems reasonable since the decision process needs time and is based on recent deviations from a desired exchange rate path.

The first strategy is called the “leaning against the wind” strategy (LAW) and may be expressed as

$$d_t^{CB} = \alpha^{CB,L} \, (LogS_{t-2} - LogS_{t-1}), \tag{7}$$

where the intervention volume in period t depends on the difference between the logarithm of the exchange rates in t–1 and t–2. The reaction coefficient $\alpha^{CB,L}$ is constant and positive. Applying this strategy, the central bank always trades against past trends.

With the second strategy, the central bank supports a target exchange rate. The TARGET strategy is formalized as

$$d_t^{CB} = \alpha^{CB,T} \, (S_{t-1}^F - S_{t-1}) / S_{t-1}, \tag{8}$$

where the intervention volume in period t depends on the relative distance between the target rate and the exchange rate in t–1. Again, the reaction coefficient $\alpha^{CB,T}$ is constant and positive. For simplicity, we assume that the target rate is equal to the fundamental exchange rate. By this strategy the central bank aims at moving the exchange rate towards its fundamental, but does not influence the value of the fundamental exchange rate itself.

To evaluate the success of an intervention two measures are considered. One is directly concerned with the variability of the exchange rate, the other one more with deviations from a target rate. The volatility measure is defined as

$$V = \frac{100}{n-1} \sum_{t=2}^{n} \left| LogS_t - LogS_{t-1} \right| , \tag{9}$$

where the volatility is measured as the average of the absolute returns. As suggested by Guillaume et al. (1997), we prefer the absolute value of the returns to the more usual squared values. Due to the nonexistens of the forth moment in the distribution of the returns, the former quantity has a greater capacity to reflect the structure in the data.

Although the volatility measure is the most common measure it should not be the only one. For instance, it is not desirable that the central bank stabilizes the exchange rate far away from its fundamental value. Therefore, we suggest an additional measure to capture the distortion in the foreign exchange market. A distortion in the sense of deviations of the exchange rate from the target rate may be defined as

$$D = \frac{100}{n} \sum_{t=1}^{n} \left| (S_t - S_t^F) / S_t^F \right| . \tag{10}$$

The distortion measures the extend to which the exchange rate fluctuates around the fundamental exchange rate. Again, absolute values are used. Clearly, if two measures are used, a trade-off may exist. This problem has to be solved by the central bank, for instance by minimizing a loss function combined out of (9) and (10).

4.3.3 Simulations

Taking into account central bank operations, market clearing requires

$$m_t d_t^C + (1 - m_t) d_t^F + d_t^{CB} = 0 \,. \tag{11}$$

Solving (11) for the exchange rate yields

$$S_t = \frac{E_t[S_{t+1}]}{1 - (m_t d_t^C + d_t^{CB}) / \alpha^F (1 - m_t)}, \tag{12}$$

where the demand of the central bank is due to its applied intervention strategy.

To understand how intervention operations affect exchange rate dynamics, we first present an example and then discuss the results more generally. Figure 4 shows an example in the time domain for 300 periods. The top part of the figure contains a simulation run without intervention. The relation between the systematic and the unsystematic demand of the chartists is roughly equal. For 1,000 observations the volatility is computed as 0.63 and the distortion as 1.38.

The middle part of figure 4 displays a time series generated with the LAW strategy. For this regime the volatility is 0.51 and the distortion 0.97. The reduction of the volatility is a consequence of the specific intervention strategy. By "leaning against the wind" the central bank reduces the momentum of the exchange rate movement. In addition, the distortion also declines.

The bottom part of figure 4 contains a simulation run for the TARGET intervention. On the one hand, the intervention operation has reduced the distortion (D=0.84). But on the other hand, the volatility has increased (V=0.81). The TARGET intervention is in an environment dominated by trend-following chartists not successful. The reason is that if the exchange rate trends toward the fundamental exchange rate, chartists and the central

central bank trade in the same direction. In such a situation, the central bank raises the momentum of the exchange rate movement.

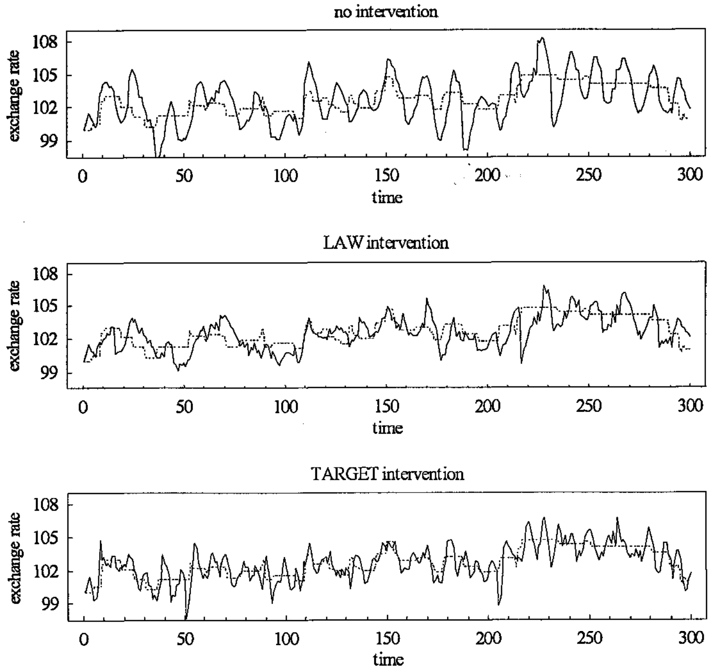

Figure 4: Example of Intervention Operations. The solid line is the exchange rate, the dashed line its fundamental, the same parameter setting as in figure 1, but $a^{C,1} = 0.75$, $a^{C,2} = 0.0045$, LAW intervention: $a^{CB,L} = 0.2$, TARGET intervention: $a^{CB,T} = 0.3$.

Note that intervention operations may also have an impact on the finer structure of the exchange rate path. Figure 4a compares the autocorrelation functions for the above time series. As can be seen, the LAW strategy clearly reduces the mean reversion tendency. In addition, if periods of intervention alternate with periods of no intervention the central bank obviously causes volatility clustering.

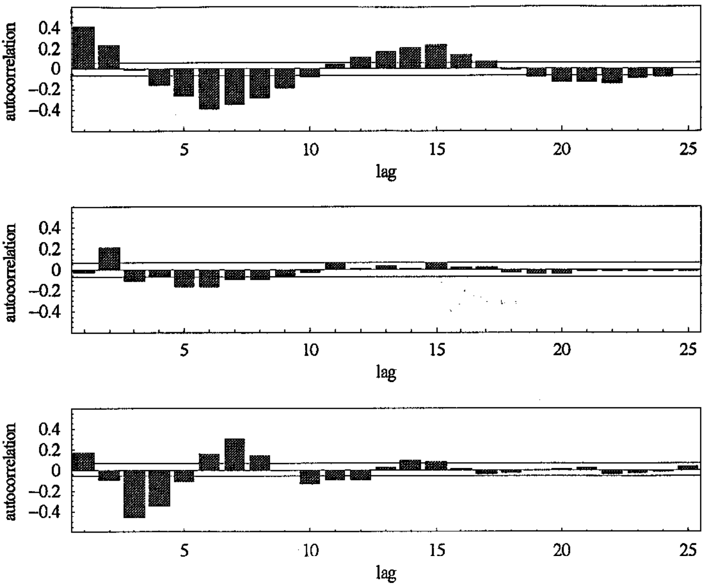

Figure 4a: Mean Reversion and Intervention Operations. The same parameter setting as in figure 4, top (middle, bottom): no (LAW, TARGET) intervention, T=1,000, 95 percent confidence intervals are plotted as $\pm 2/\sqrt{T}$ (assumption of white noise).

Figure 5 shows the simulation results in more detail. Each part of figure 5 is constructed as follows. On the one axis the intervention volume, that is the coefficient a^{CB}, is increased in 20 discrete steps starting from zero. On the other axis the noise level $a^{C,2}$ is increased. To hold the volatility constant, the influence of the first part of the technical trading rule, as given by $a^{C,1}$, is lowered. In this way the volatility is fixed around 0.6 (no intervention!).[9] The initial noise level is zero and increased in 20 discrete steps. The maximum (medium) noise level results in an unsystematic trading volume that is roughly 2/3 (1/2) of the total. Thus, the figure displays regimes where the chartists trade rather unsystematically as well as regimes where the trading positions are correlated. The volatility and the distortion are calculated on the basis of 1,000 observations.

[9] For major currencies the volatility ranges between 0.4 and 0.5 (compare section 2.1.2.2). In Westerhoff (2001), the volatility is set somewhat lower than observed empirically (V=0.3). The impact of central bank interventions appears to be similar in calm periods (compare also appendix 4.A.1)

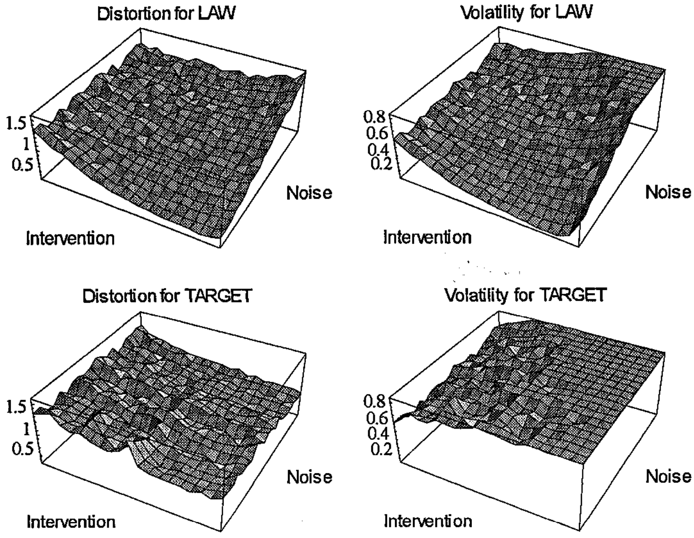

Figure 5: Intervention Results. Distortion and volatility are calculated on the basis of 1,000 observations, the same parameter setting as in figure 1, the intervention level is increased in 20 steps from 0 to 0.475, the noise level in 20 steps from 0 to 0.0076, β^2 is identified so that for $a^{C,1}=1$ and $a^{C,2}=0$ the volatility is approximately 0.6, to hold the volatility constant while increasing the noise level, $a^{C,1}$ is appropriately reduced.

As can be seen, the LAW strategy is in a trend-following environment very effective. The volatility can be reduced depending on the noise level. The reason is that the central bank, by "leaning against the wind", reduces the momentum of the exchange rate movement and therefore the power of the trading signals. When the noise level rises, the effectiveness declines. Note that whenever the direction of the trend changes, the central bank, due to their decision lag, intervenes in the wrong direction. Hence, if the behavior of the chartists becomes less trend-following, the success of LAW intervention diminishes. As a byproduct, the distortion declines also, especially when the noise level is low.

The TARGET intervention does not reduce the volatility. The main reason is that if the exchange rate moves in the direction of the fundamental exchange rate, chartists and the central bank typically trade in the same direction. This strategy is able to reduce the distortion, yet at the cost of higher volatility. Therefore, the LAW strategy seems to be preferable, especially if the central bank is sure that the exchange rate trend prevails.

4.3.4 Periods of High Uncertainty

If the environment is extremely uncertain, the agents allow themselves to be guided by past values of the exchange rate when forming new expectations. These function as "anchors" in the individual judgement of the future exchange rate. This phenomenon is called anchoring heuristics and well documented in the psychological literature (Tversky and Kahneman 1974). In such periods the expectation formation process, with respect to the exchange rate, is not only regressive but also anchored to the last few observations of the exchange rate. Assuming $E_t[S_{t+1}]=\gamma S^F_{t-1}+(1-\gamma)(S_{t-1}+S_{t-2})/2$, equation (2) modifies according to

$$d_t^F = \alpha^F (\gamma S_{t-1}^F + (1-\gamma)(S_{t-1} + S_{t-2})/2 - S_t)/S_t, \qquad (13)$$

where the fundamentalists now use the exchange rate in t–1 and t–2 as an orientation for expectation formation.

The top of figure 6 displays the implications of anchor expectations for the exchange rate behavior. The main difference is that after a stronger outlier occurs, for instance triggered by a shock, the exchange rate stays near the new exchange rate for a while before reversion sets in. This is a consequence of the anchoring behavior, since the agents stick to the past when forming expectations. For the first 1,000 observations the volatility is 0.60 and the distortion is 0.91.

What are the implications for central bank interventions in such an environment? The middle part of figure 6 shows a simulation run for LAW interventions. Through the intervention operation the volatility is increased. The distortion is not influenced (V=0.72, D=0.90). In this regime, the LAW intervention is not successful since the market is not trending. In contrast to regressive expectations, the anchor heuristic together with the high uncertainty often leads to reversals of the exchange rate movement with the result that the central bank intervenes in the wrong direction. Moreover, in some cases the adjustment of the exchange rate towards its fundamental is slowed down. Instead of smoothing the exchange rate path, the exchange rate fluctuates more strongly and in a greater distance to its fundamental value. Hence, the distortion is not reduced.

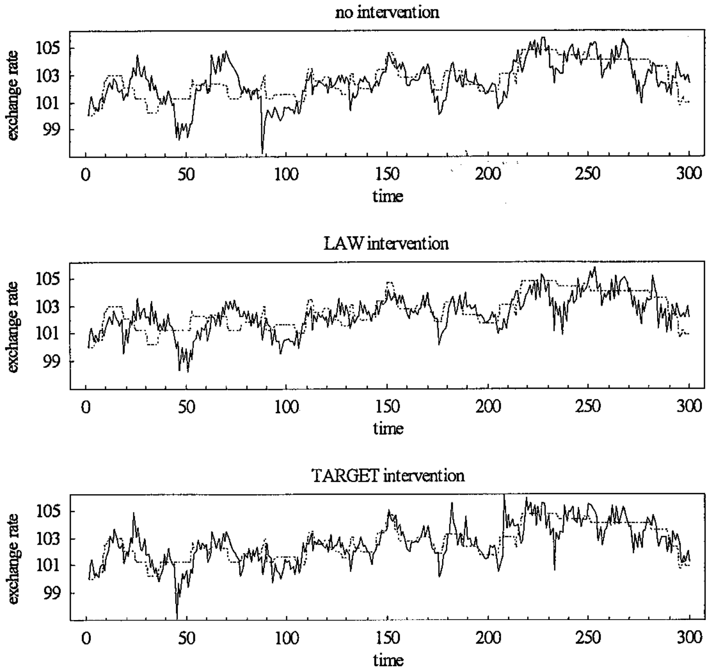

Figure 6: Example of Intervention Operations for Anchor Expectations. The solid line is the exchange rate, the dashed line its fundamental, $S_1^F=100$, $S_1=100$, $S_2=101$, $S_3=101.5$, $a^{C,1}=0.24$, $a^{C,2}=0.003$, $a^F=1$, $\beta^1=0.1$, $\beta^2=9.5$, $\gamma=0.2$, $\delta \sim N(0,1)$, $prob(p=1)=0.2$, $\varepsilon \sim N(0, 0.0075)$, T=300, LAW intervention: $a^{CB,L}=0.1$, TARGET intervention: $a^{CB,T}=0.1$.

The bottom displays the impact of the TARGET intervention. The results are mixed. The volatility has increased to 0.70, but the distortion has dropped to 0.64. In periods of high uncertainty the central bank seems to be able to improve the distortion at the cost of a higher volatility. Since the exchange rate is always pushed towards its fundamental value, the volatility automatically rises.

Figure 7 shows the consequences of the intervention operations in more detail. The figure is constructed in the same way as before. In general, the LAW operations have not the power to reduce the volatility. Only if the noise level is low, the distortion might be reduced. But it seems natural that in periods of high uncertainty chartists do not rely on trend-following trading rules. In such times, their behavior is better described as unsystematic.

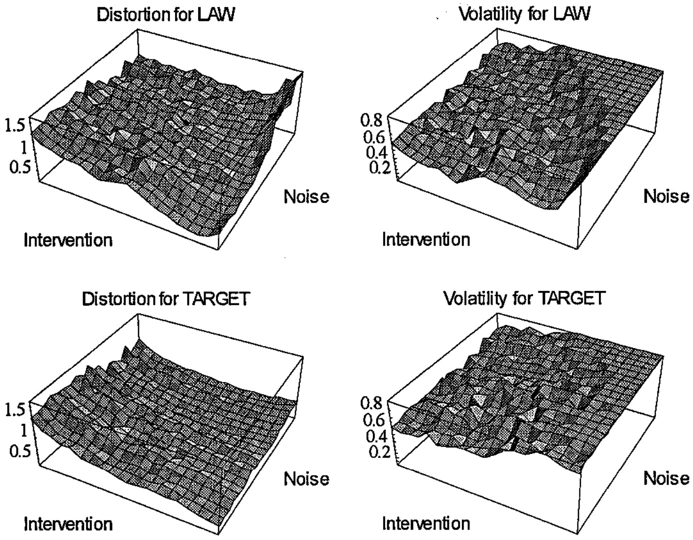

Figure 7: Intervention Results for Anchor Expectations. Distortion and volatility are calculated on the basis of 1,000 observations, the same parameter setting as in figure 6, the intervention level is increased in 20 steps from 0 to 0.475, the noise level in 20 steps from 0 to 0.0038, β^2 is identified so that for $a^{C,1}=1$ and $a^{C,2}=0$ the volatility is approximately 0.6, to hold the volatility constant while increasing the noise level, $a^{C,1}$ is appropriately reduced.

Although TARGET operations increase the volatility they allow to reduce the distortion. By using this volatility-enhancing strategy the central bank stabilizes the exchange rate near its fundamental value. In a broader sense, the central bank takes over the role of the fundamentalists.[10]

[10] Compare appendix 4.A.2 for a low volatility regime.

Note that our model replicates the empirical findings of Hung (1997). On the one hand, central bank interventions are able to reduce the volatility, but on the other hand the opposite effect is also observable. An increase in the volatility may be the price for less distorted markets. However, volatility decreasing interventions seem to be the more used strategy.[11, 12]

4.3.5 Support of Fundamentalism

So far, the means of the central bank may be evaluated rather pessimistically. Only if the chartists behave trend-following, LAW can be successful. The TARGET strategy may decrease the distortion, but leads to a higher volatility. In addition, the costs of the intervention operations are unclear. Another problem could be that for a successful operation the intervention volume is unreasonably high. Finally, if it is not possible to sterilize these transactions other markets may be disturbed.

Besides direct intervention operations, the central bank may wish to control the dynamics indirectly. In our model this can be reached by encouraging fundamental trading. Remember that the influence of the fundamentalists is controlled by two parameters: β^1 reflects the ground proportion of fundamental traders and β^2 the popularity of fundamental trading rules. Now, if the central bank provides better information about the fundamental exchange rate, β^1 and β^2 might increase. Figure 8 shows that such a support of the fundamentalists reduces both the distortion and the volatility. This holds independently of the noise level.

---

[11] To check the robustness of the intervention outcome, we repeated the simulations with other functional and numerical specifications. For instance, we used in (1) a double crossover method (instead of the moving average rule), or in (4) a quadratic weighting scheme (instead of the square root). The qualitative results remain stable under such modifications.

[12] In reality, central banks intervene, of course, less frequently. Thus, we modified (7) and (8) so that an intervention is only triggered if $|LogS_{t-1} - LogS_{t-2}|$ or if $|(S^F_{t-1} - S_{t-1})/S_{t-1}|$ exceeded a certain threshold. Varying these thresholds, we found that the intervention operation works the best, if the central bank permanently intervenes as specified by (7) and (8).

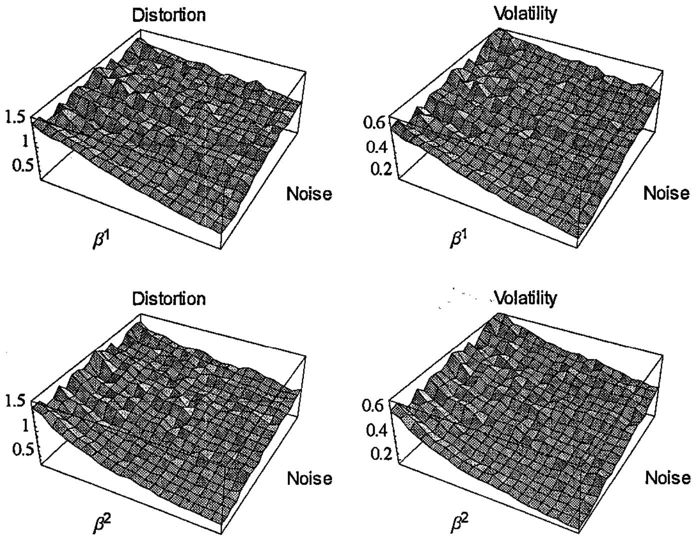

Figure 8: Support of Fundamentalism. Distortion and volatility are calculated on the basis of 1,000 observations, the same parameter setting as in figure 1, but β^1 (β^2) is increased in 20 steps from 0.1 to 0.5 (from 22 to 52), the noise level in 20 steps from 0 to 0.008, to fix the volatility around 0.6 while increasing the noise level, $a^{C,1}$ is reduced.

4.4 Conclusions

The aim of this chapter is, first, to develop a better understanding of the driving forces of exchange rate dynamics, and second, to study whether typical intervention strategies are able to reduce the high volatility. Guided by empirical observations, the focus of our analysis is on the speculative behavior of the traders. To conclude, we have identified four main forces responsible for the complex dynamics.

- First, technical trading rules typically destabilize the market. Especially when the demand of chartists is correlated, i.e. if they trade systematically into one direction, stronger trends in the exchange rate path are observed.
- Second, fundamental trading rules typically stabilize the market. But when the agents are uncertain about the fundamental value of the exchange rate, these strategies may also contribute to a distortion in the foreign exchange market.

- Third, another source is, of course, the news arrival process. Although the news arrival process is the classical argument explaining foreign exchange dynamics it is only one factor among others. Our model shares even for a low probability of fundamental shocks some important stylized facts of the empirical data: a high variability of the exchange rates, fat tails for returns, and weak evidence of mean reversion.
- Fourth, central bank interventions also have an impact on the dynamics. By "leaning against the wind", the autocorrelation of the returns may be reduced. In contrast, other financial markets exhibit a stronger tendency of mean reversion although the speculative investment positions of the agents are derived in a similar way. Moreover, if periods of intervention alternate with periods of no intervention, the central bank induces a volatility clustering.

Depending on what drives the dynamics, the central bank may be able to stabilize the market by intervention.

- If the investment positions of the chartists are correlated, a "leaning against the wind" strategy is able to reduce the volatility.
- If the market participants are uncertain about the fundamental value of the exchange rate, the central bank has the opportunity to reduce the distortion by supporting its target exchange rate.

Apart from direct interventions, the central bank may also encourage the fundamentalists to take more risk by providing better information about the fundamental value of the exchange rate.

4.A Appendix

4.A.1 A Low Volatility Regime and Regressive Expectations

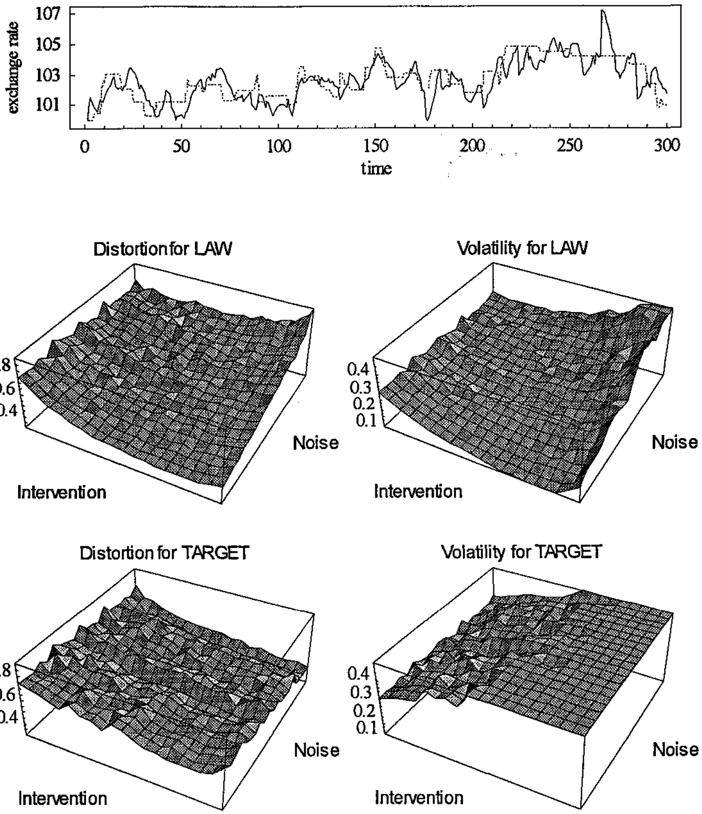

Figure 8: Low Volatility and Regressive Expectations. In the top the exchange rate in the time domain, S_1^F=100, S_1=100, S_2=101, S_3=101.5, $a^{C,1}$=0.4, $a^{C,2}$=0.003, a^F=1, β^1=0.1, β^2=30, γ=0.2, δ~N(0,1), $prob$(p=1)=0.2, $\varepsilon \sim N(0, 0.0075)$, T=300. In the bottom the distortion and the volatility, calculated on the basis of 1,000 observations using the same parameter setting, the intervention level is increased in 20 steps from 0 to 0.475, the noise level in 20 steps from 0 to 0.00285, β^2 is identified so that for $a^{C,1}$=1 and $a^{C,2}$=0 the volatility is approximately 0.3, to hold the volatility constant while increasing the noise level, $a^{C,1}$ is appropriately reduced.

4.A.2 A Low Volatility Regime and Anchor Expectations

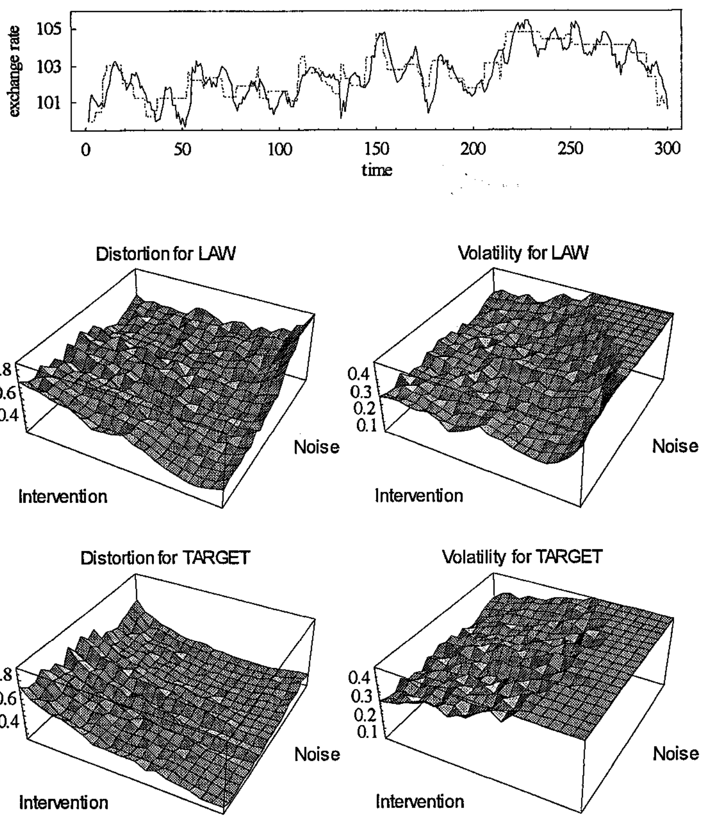

Figure 9: Low Volatility and Anchor Expectations. In the top the exchange rate in the time domain, $S_1^F=100$, $S_1=100$, $S_2=101$, $S_3=101.5$, $a^{C,1}=0.6$, $a^{C,2}=0.0012$, $a^F=1$, $\beta^1=0.1$, $\beta^2=15$, $\gamma=0.2$, $\delta \sim N(0,1)$, $prob(p=1)=0.2$, $\varepsilon \sim N(0, 0.0075)$, T=300. In the bottom the distortion and the volatility, calculated on the basis of 1,000 observations using the same parameter setting, the intervention level is increased in 20 steps from 0 to 0.475, the noise level in 20 steps from 0 to 0.0019, β^2 is identified so that for $a^{C,1}=1$ and $a^{C,2}=0$ the volatility is approximately 0.3, to hold the volatility constant while increasing the noise level, $a^{C,1}$ is appropriately reduced.

5 Expectations Driven Distortions in the Foreign Exchange Market

Abstract

This chapter* explores the phenomenon of lasting deviations of the exchange rate from its fundamental value in the foreign exchange market. Motivated by empirical observations a chartists-fundamentalists model is developed in which boundedly rational agents repeatedly choose between technical and fundamental trading rules to determine their speculative investment positions. Crucial for the dynamics is how the traders perceive the fundamental exchange rate. This perception process is based on psychological evidence. Simulations give rise to bubbles but simultaneously display quite realistic exchange rate dynamics (unit roots in the exchange rates, fat tails for returns, and volatility clustering).

Keywords

exchange rate theory, technical and fundamental trading rules,
expectations and learning, market efficiency

JEL Classification

D84, F31, G14

* Presented at the 7th International Conference of the Society for Computational Economics on Computing in Economics and Finance, Yale University, New Haven, June 2001 and at the 2nd CeNDEF Workshop on Economic Dynamics, Amsterdam, January 2001. I thank the participants for their helpful discussions, especially Cars Hommes and Allan Timmermann.

5.1 Introduction

This chapter aims at explaining the phenomenon of distortions in the foreign exchange market. Distortions in the sense of (lasting) deviations of the exchange rate from its fundamental value are a sign of market inefficiency. One well-known example for this is the bubble path of the US dollar in the eighties. In January 1980, the mark-dollar exchange rate was around DM 1.70. In the next five years, the exchange rate increased over 100 percent. Its height of DM 3.46 was reached in February 1985. Afterwards, the exchange rate depreciated sharply by over 50 percent. At the end of 1987, the exchange rate dropped below DM 1.60.

More specifically, we want to develop a model which gives rise to bubble processes but yields realistic exchange rate movements at the same time, namely unit roots in exchange rates, fat tails for returns, and volatility clustering. Further aims include modeling the perception of the fundamental exchange rate more explicitly than usual on the grounds of psychological evidence and finally separating some of the forces responsible for the distortions.

Related (behavioral finance) research has already produced several papers which replicate some of the stylized facts of financial markets (de Grauwe et al. 1993, Lux 1997, Brock and Hommes 1997 and 1998, Farmer 2000). Moreover, the approach of Kirman (1991) is able to produce short-lived bubbles, whereas Frankel and Froot (1986) find an explanation for lasting bubbles. In their model, the bubble process is fed by a smoothly and slowly declining influence of the fundamentalists. The bubble path is at its turning point when the impact of the fundamentalists is at its minimum. In addition to the absence of typical time series properties, such a weighting scheme does not always hold.

The starting point for our investigation is a simple chartists-fundamentalists model. Motivated by empirical observations, a model is developed where boundedly rational market participants choose between a technical and a fundamental trading rule to determine their speculative investment positions. This decision is repeated at the beginning of each new trading period. If one subscribes to the strong assumption that the agents are able to determine the true fundamental value of the exchange rate, then the exchange rate fluctuates in a complex way around its fundamental value. The computed time series mimic the behavior of major currencies quite closely. Besides some excess volatility, the foreign exchange market seems to be efficient.

However, in this chapter we try to go one step further and model the perception of the fundamental exchange rate more realistically. While the agents follow the news arrival process closely, mistakes in information processing occur. These mistakes are propagated over time since the agents tend to stick to their previously perceived fundamental value (anchor heuristic). If the agents believe that the exchange rate itself contains relevant information, they incorporate it into their "anchor" so that the exchange rate becomes even more disconnected from its true fundamental. Nevertheless, in the long-run the agents react to macroeconomic imbalances and adjust their perception. Through this learning procedure, the distortion eventually is corrected and some long-term mean-reversion sets in. Note that the results are not the outcome of strange exchange rate movements. On the contrary: the generated time series share some basic stylized facts with the empirical data.

This chapter is organized as follows. Section 5.2 presents a chartists-fundamentalists framework and discusses its time series properties. Adopting heuristics from the psychological literature, section 5.3 modifies the perception of the fundamental exchange rate. This extension gives rise to bubble processes. Section 5.4 offers some conclusions.

5.2 A Simple Chartists-Fundamentalists Model

First of all, the agents considered in this chapter are not fully rational. Therefore, their behavior is not derived out of a well-defined utility maximization problem. Instead, we provide an empirical micro-foundation for the trading activity of the agents by using, for instance, observations from the market micro-structure or psychological evidence. Our point of reference is that the agents are boundedly rational. According to Simon (1955), this term recognizes the cognitive limitations of a decision-maker with respect to both knowledge (including the relevant theory and necessary information) and computational capacity.

Psychologists like Tversky and Kahneman (1974) have produced a huge amount of experimental material which impressively demonstrates that people tend to rely on a limited number of heuristic principles. The complex tasks of assessing probabilities or predicting values are reduced to simpler judgmental operations hereby. It should be clear that heuristics are a kind of constraint. They prevent the agents from making an infinite number of useless inferences, but they also prevent the agents from making a much smaller number of useful inferences (Holyoak and Nisbett 1988). In general, heuristics are quite useful, but sometimes they lead to severe and systematic errors.

Noteworthy examples for such failures are given in survey studies on expectation formation (Ito 1990, Takagi 1991). Most importantly, market participants are heterogeneous, e.g. there exist individual peculiarities like wishful thinking which clearly violate the rational expectation hypothesis. Furthermore, the agents appear to adhere to destabilizing bandwagon expectations in the short-run, but display a stabilizing mean-reversion in the long-run.

For Heiner (1983), who argues more in economic terms, the limits of maximizing are the origin of predictable behavior. Observed regularities of behavior are the outcome of behavioral rules. For example, the specific complexity of the foreign exchange market leads for every agent to a gap between his competence in making an optimizing decision and the actual difficulty involved with this decision. The wider the gap, the more likely agents will follow a rule-governed strategy. This makes their behavior predictable.

Two observations are crucial for our model. In recent years, the daily foreign exchange turn-over has increased sharply. More and more, the trading volume reflects very short-term transactions, indicating a highly speculative component (BIS 1999). Surprisingly, when the speculators determine their investment position, rather simple trading rules are applied. Survey studies (Taylor and Allen 1992, Menkhoff 1997, Lui and Mole 1998) unanimously confirm that most of the professional foreign exchange traders rely on both technical and fundamental trading rules.

Thus, our model lives from the facts that the foreign exchange market is dominated by speculative activity and that the traders are familiar with both technical and fundamental trading methods. The crucial idea of the model may be summarized as follows. At the beginning of every trading period, the traders choose a specific trading rule to determine their speculative investment position. The traders have the choice between technical and fundamental trading rules. Their selection depends on expected future performance possibilities.

5.2.1 Setup of the Model

Technical analysis is a trading method that attempts to identify trends and reversals of trends by inferring future price movements from those of the recent past (see Murphy 1999 for a popular tutorial of technical analysis). Trading signals are spotted by applying both graphical (charts) and statistical tools. To cover these trading methods, the technical demand of the agents for period t is divided into two components

$$d_t^C = a^C \{a^{C,1}(0.6(LogS_{t-1} - LogS_{t-2}) + 0.4(LogS_{t-2} - LogS_{t-3})) + a^{C,2}\, \delta_{t-1}\}. \quad (1)$$

The first one reflects the typical behavior of technical traders, where the trading signals are triggered by a simple moving average rule. In general, chartists buy (sell) foreign currency if the exchange rate S rises (declines). The second one represents additional random demand to allow for more complicated behavior. The stochastic variable δ is assumed to be normally distributed with mean zero and constant variance. The positive reaction coefficients α^C, $\alpha^{C,1}$, and $\alpha^{C,2}$ calibrate the total demand of (1) and the relation between the systematic and unsystematic component. Finally, note that in (1) a market order for period t is generated in response to past price changes. Such a lag structure is typical for technical trading rules since only past movements of the exchange rates are exploitable by these rules.

In contrast to technical trading rules, fundamental trading rules aim to gain from differences between the exchange rate and its fundamental value. Fundamentalists believe that the exchange rate will converge towards its equilibrium value in the future (Moosa 2000). The demand of the fundamentalists is formalized as

$$d_t^F = a^F \{(E_t[S_{t+1}] - S_t)/S_t\}, \quad (2)$$

where α^F is a positive reaction coefficient. The fundamental trading rule delivers a buy (sell) signal, if the expected future exchange rate is above (below) the spot rate. The amount of the demand depends on the relative distance between the expected rate and the spot rate.[1]

These expectations are modeled in a classical regressive manner. If the exchange rate deviates from its perceived fundamental value S^{FP}, a turn-back is expected. Thus

$$E_t[S_{t+1}] = \gamma_{t-1} S_{t-1}^{FP} + (1-\gamma_{t-1}) S_{t-1}, \tag{3}$$

where γ stands for the expected adjustment speed of the exchange rate towards its fundamental value. For instance, if γ is 0.25, the agents expect an adjustment of 25 percent. However, γ is not constant but drawn from a Uniform distribution. Since the expectation formation process for the trading period t has to be made in advance, the last available data is from period t−1.

Decisive to this chapter is the way in which the agents form their perception of the true fundamental exchange rate. This will be discussed in the next section. Here, we will simply assume that the agents are able to determine on average the true fundamental exchange rate S^F, although they do make some mistakes in every period

$$LogS_t^{FP} = LogS_t^F + \lambda_t, \tag{4}$$

where the mistake λ is normally distributed with mean zero and constant variance.

The development of the fundamental exchange rate is due to the news arrival process. The fundamental value follows a jump process. Its logarithm is given as

$$LogS_t^F = LogS_{t-1}^F + \eta_t, \tag{5}$$

where the news η is identically and independently distributed according to a Normal distribution with mean zero and constant variance. However, news do not occur in every trading period. In such a case η is zero.

The decision for a trading rule has to be made before the trading starts. The selection depends on expected future performance possibilities. Fundamentalism, compared to

[1] Due to the time structure of the model, the fundamentalists function as market makers. Technical traders derive their orders from past price movements, the market clearing is established by the fundamentalists. They adjust the price in order to absorb the total demand.

chartism, becomes more popular the wider the spot rate deviates from its expected future rate. The weight of the chartists is defined as

$$m_t = \frac{1}{1+\beta^1+\beta_{t-1}^2\sqrt{|E_t[S_{t+1}]-S_{t-1}|/S_{t-1}}}, \tag{6}$$

and the weight of the fundamentalists as $(1-m_t)$. The coefficient β^1 represents the basic influence of the fundamentalists, i.e. $(1-1/(1+\beta^1))$ is the minimum fraction of agents who are always fundamentalists. Nevertheless, most of the traders adjust their trading strategy with respect to the relevant market conditions. As indicated by (6), the weight of the fundamentalists increases, though at a declining rate, as the relative distance between $E_t[S_{t+1}]$ and S_{t-1} rises. In such a situation, more and more of the speculators come to the conclusion that the exchange rate is mispriced so that fundamental analysis is preferable to technical analysis. The time dependent coefficient β_t^2 reflects the popularity of fundamental analysis and stems from a Uniform distribution.

Demand from the international trade and the risk management of the firms play, compared to speculative transactions, a minor but still significant role. The liquidity needs of the firms for period t is given as

$$d_t^{IT} = \alpha^{IT}\{\alpha^{IT,1}\chi_{t-1}+\alpha^{IT,2}(S_{t-1}-S_{t-1}^F)/S_{t-1}^F\}, \tag{7}$$

where α^{IT}, $\alpha^{IT,1}$, and $\alpha^{IT,2}$ are reaction coefficients. The first source of the demand χ is normally distributed (with mean zero and constant variance), whereas the second source reflects the usual current account relationship. For example, if the exchange rate is higher than its fundamental, then exports exceed imports. Since the focus is on rather short time periods, say daily time intervals, the sign of the total demand of the firms is not a priori clear. Even a medium current account imbalance may be overcompensated by the random component.

The market clearing condition is given as the sum over all transactions in period t

$$m_t d_t^C + (1-m_t)d_t^F + d_t^{IT} = 0. \tag{8}$$

Solving (8) for the exchange rate yields the solution of the model

$$S_t = \frac{E_t[S_{t+1}]}{1-(m_t d_t^C + d_t^{IT})/(1-m_t)}, \tag{9}$$

which is a four-dimensional stochastic difference equation system. As can be inferred from (9), larger price reactions occur whenever a low proportion of fundamentalists is confronted with a huge demand of chartists.[2] Since (9) precludes closed analysis, simulations are performed to demonstrate that the underlying structure gives rise to complex exchange rate motion as it is observed empirically.[3]

5.2.2 Calibration

Before we discuss the simulation results, the model has to be calibrated. Unfortunately, only some of the parameters are empirically observable. The parameter setting is displayed in table 1. With the help of the reaction coefficients, the demand of the traders and the firms is controlled. The relation between the systematic and unsystematic demand of the technical trading rule is around 1 to 6. In comparison, the demand of the firms is about one third of the total (BIS 1999).[4] Survey studies (Taylor and Allen 1992) show that roughly 15 percent of the market participants always rely on fundamental analysis, i.e. $\beta^1=0.177$. The coefficient β^2 is chosen so that the variance of the exchange rate time series matches the variance of daily exchange rate movements. Note that the popularity of being a fundamentalist is not constant, but changes randomly from time to time. This might be justified by changing economic conditions which result in periods of higher and lower uncertainty about the fundamental exchange rate. The development of the fundamental exchange rate takes into account that fundamental shocks do not occur in every period. In contrast to the news arrival process, the misperception of the fundamental exchange rate is rather small. Finally, the agents expect an adjustment of the exchange rate towards its fundamental between 0 and 50 percent. On average, γ changes every 6 trading periods.

[2] Because there is always a ground proportion of fundamentalists in the market, the system does not result in unrealistically strong exchange rate movements (compare table 2).

[3] The mechanics of the model are best described by a stretching and folding of the exchange rate around its perceived fundamental value. In the absence of noise, such a mechanism has the potential to generate chaotic motion. A deterministic skeleton of (9) is discussed in chapter 3.

[4] The firms' demand is mainly random; only 0.3 percent of total transactions are due to the current account relationship. However, this changes in section 5.3. Lasting bubbles induce an increase of the current account demand up to 1.7 percent of total transactions, which is consistent with BIS estimates.

description of parameter	symbol	value
reaction coefficient Eq. (1)	α^{C}	1
reaction coefficient Eq. (1)	$\alpha^{C,1}$	0.225
reaction coefficient Eq. (1)	$\alpha^{C,2}$	0.008
reaction coefficient Eq. (2)	α^{F}	1
reaction coefficient Eq. (7)	α^{IT}	1
reaction coefficient Eq. (7)	$\alpha^{IT,1}$	0.0014
reaction coefficient Eq. (7)	$\alpha^{IT,2}$	0.001
random demand of traders Eq. (1)	δ	N(0, 1)
random demand of firms Eq. (7)	χ	N(0, 1)
basic level of fundamentalists Eq. (6)	β^{1}	0.177
popularity of fundamentalists Eq. (6)	β^{2}	Prob(β_t^2~U(8, 100))=2/52, else $\beta_t^2=\beta_{t-1}^2$
distribution of news Eq. (5)	η	Prob (η_t ~ N(0, 0.0075)) = 1/5, else $\eta_t = 0$
misperception Eq. (4)	λ	N(0, 0.001)
expected adjustment Eq. (3)	γ	Prob (γ_t~ U(0, 0.5)) = 1/6, else $\gamma_t = \gamma_{t-1}$

Table 1: Basic Parameter Setting for the Simulations.

5.2.3 Snapshot of the Dynamics

This section demonstrates that simulations of the model are able to replicate three important univariate stylized facts of exchange rate fluctuations: (i) unit roots in the exchange rates, (ii) fat tails for returns, and (iii) volatility clustering (for general surveys, compare Guillaume et al. 1997, Pagan 1996, and de Vries 1994).

Figure 1 gives a first impression of the dynamics. The top contains the exchange rate (solid line) and the development of its fundamental (dashed line) for 300 periods, the bottom displays the corresponding weights of the chartists. Even a low probability of fundamental shocks suffices to generate realistic exchange rate movements where the spot rate circles around its fundamental in a complex fashion. The volatility of the exchange rate is clearly higher than its fundamental value. The weight of the technical traders is mostly concentrated in the range from 30 to 60 percent with some peaks going down to 15 or up to 75 percent. In other words, the agents rely on both kinds of investment strategies.

Figure 1 may help to explain the fuzzy relationship between news and exchange rate movements. Goodhart (1988) reports both systematic underreaction and overreaction to news. Even large price movements unrelated to news are apparent. Visual inspection of figure 1 reveals similar findings. Our model suggests that the dynamics are partially caused endogenously through the interactions between the traders. For instance, after a shock has hit the market, the exchange rate reaction may be amplified and prolonged by the positive feedback trading of the technicians. Black (1986) has called this behavior noise trading. According to Black, noise, which can be described as a large number of small events, is essential to the existence of liquid markets. He argues that a person who wants to trade needs another person with opposite beliefs. To explain the high trading volume, it is not reasonable to assume that differences in beliefs are merely the outcome of different information. In this model, the traders may even generate their own trading signals in periods with no new information at all.

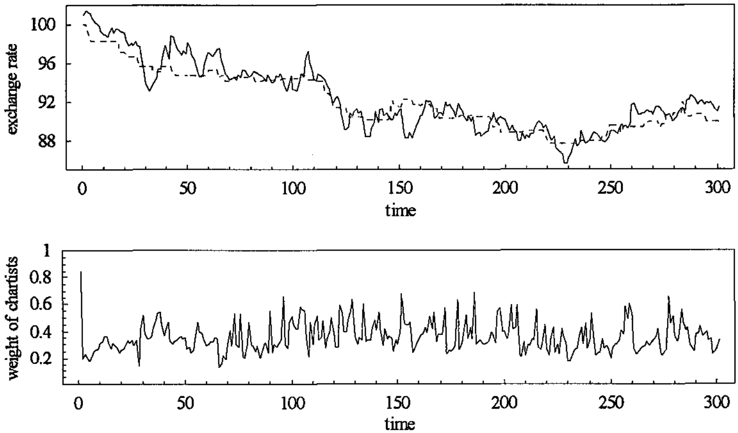

Figure 1: Exchange Rate Dynamics and Weight of Chartists. The solid line is the exchange rate, the dashed line its fundamental. Parameters are as in table 1.

One stylized fact of the empirical literature is that exchange rate time series display unit roots (Goodhart et al. 1993). We have tested the null hypothesis that any shock to the exchange rate is permanent against the alternative hypothesis that a shock is only temporary with the augmented Dickey-Fuller test (Dickey and Fuller 1979, 1981). Using a four-lag specification without intercept delivers

$$\Delta S_t = a_0 S_{t-1} + a_1 \Delta S_{t-1} + a_2 \Delta S_{t-2} + a_3 \Delta S_{t-3} + a_4 \Delta S_{t-4}. \qquad (10)$$

Since a_0 is not significant for the computed time series the null is not rejectable. This holds for various lag settings of (10) and values of the coefficients (compare also table

4). Often, this result is interpreted as evidence for a random walk behavior of the exchange rate. Looking at figure 1, the reason for the strong support of the unit root hypothesis for our model becomes obvious. The exchange rate fluctuates around its fundamental, which itself follows a stochastic process. Since the exchange rate path never gets strongly disconnected from its fundamental value for a longer period, it also appears to be a random process. Inspecting the first 10,000 periods, one finds that the difference between both time series stays mainly in the region of ±5 percent. In the most extreme cases, the exchange rate diverges up to 15 percent from its fundamental. However, such deviations are very short-lived. Exactly speaking, the exchange rate and its fundamental are highly cointegrated. Relaxing this property is the aim of section 5.3.

Another stylized fact states that the distribution of the returns has fat tails (Guillaume et al. 1997). Returns are defined as Log-price changes (r_t=LogS_t–LogS_{t-1}). Relative to a Normal distribution with identical variance, one finds a stronger concentration around the mean, thinner shoulders, and more probability mass in the tails of the distribution. Figure 2 compares the distribution of the returns and its scaling behavior for the simulated returns and for normally distributed returns. Besides visual examination, the estimates of the kurtosis reveal fat tails. Table 2 summarizes some descriptive statistics of a simulation run over 10,000 periods. Since the kurtosis is higher than 3 (the theoretical value of a Normal distribution), the computed time series possesses fat tails. Note that the largest exchange rate movements are not unrealistically high, but they match empirical observations. As already mentioned, the model is calibrated so that the variance fits with daily exchange rate movements of major currencies. Compared to the variance of r_t^F=LogS_t^F–LogS_{t-1}^F, the variance of the exchange rate returns is six times higher. This clearly reflects strong excess volatility.

min	median	max	variance	skewness	kurtosis
6.49 %	0.00 %	6.63 %	0.000061	0.1136	9.73

Table 2: Descriptive Statistics of Returns. Parameters are as in table 1, 10,000 observations.

Fat tails may also be detected by determining the tail index. The tail index α, given as F(|return|>x)≈cx$^{-\alpha}$, 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. Figure 2 illustrates that the distribution of the returns roughly follows a power law. A regression on the largest 30 percent of the observations delivers a significant tail index of 3.64, which is consistent with results obtained from empirical data (Lux and Ausloos 2000). But what causes fat tails? The model does not only produce stronger outliers as a result of strong random demand shocks, but also if a medium demand has to be absorbed by a low weight of the fundamentalists. The tail index for a Normal distribution (see figure 2) is clearly higher.

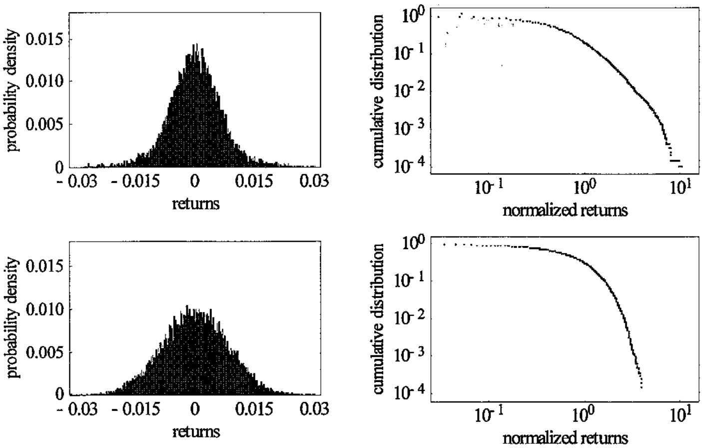

Figure 2: Distribution of Returns and Scaling Behavior. The top shows the distribution of the returns and the scaling behavior of the cumulative distribution of the positive and negative tails for normalized Log-returns (parameters are as in table 1, 20,000 observations), the bottom the same but for a Normal distribution with identical variance.

The third stylized fact highlights that periods of low volatility alternate with periods of high volatility (Mandelbrot 1963). Although almost no autocorrelation exists for raw returns, a different picture emerges if one uses absolute returns as a volatility measure (figure 3).[5] The autocorrelation for absolute returns is clearly significant and slowly decaying.[6] The reason for the short-term volatility clustering lies in the feedback trading

[5] Figure 3 reveals a weak tendency of short-run mean-reversion which is also observable in the empirical data (Cutler et al. 1990).

[6] The autocorrelation for squared returns is lower but still significant (not displayed). This hints at a multi-scaling behavior (Lux and Ausloos 2000). Calculating Hurst coefficients for $|r_t|^i$ for i=1−10 supports this conjecture. The coefficients drop continuously from 0.71 to 0.58.

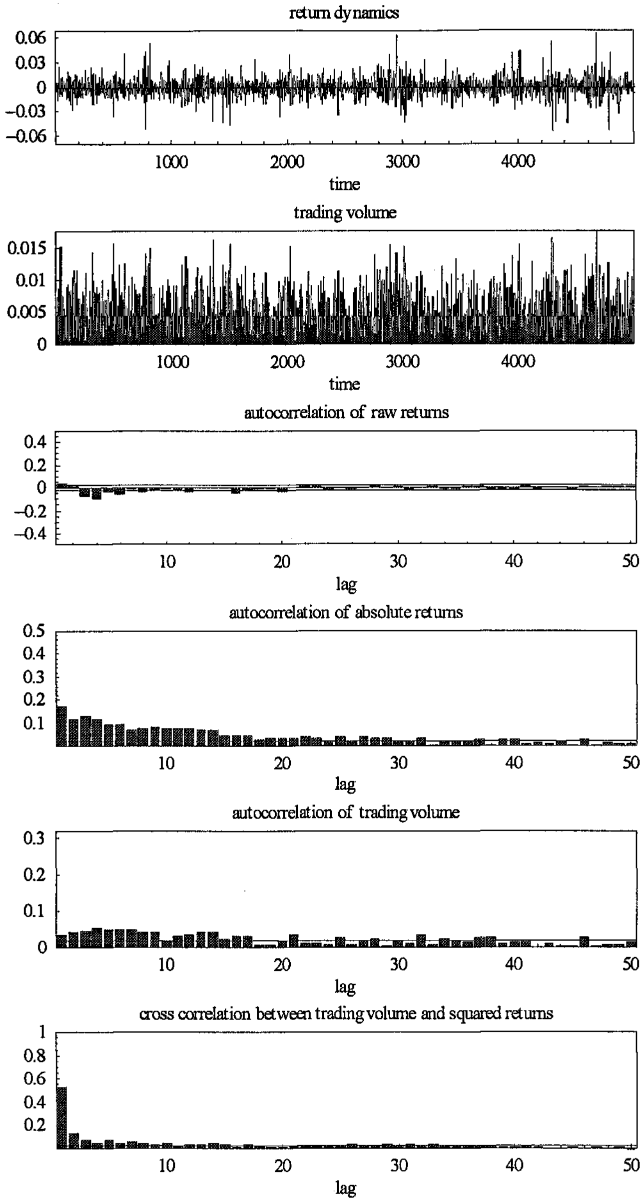

Figure 3: Volatility Clustering. Parameters are as in table 1, T=5,000 observations. 95 percent confidence intervals are plotted as $\pm 2/\sqrt{T}$ (assumption of white noise).

of the agents. A strong exchange rate movement in period t indicates a strong trading signal for period t+1. The long-run autocorrelation stems from different degrees of the popularity of the fundamental trading rules.[7] Note that the trading volume also tends to cluster. However, the cross-correlation of volatility is positive only with current volumes but almost zero for past and future volumes (Brock and LeBaron 1996).

So far, the model is able to replicate some important stylized facts. But the strong connection between the exchange rate and its fundamental seems not to be very realistic. For instance, in this case the path of the US dollar in the eighties would be fundamentally justified. In addition, it would imply that the foreign exchange market is more or less efficient (besides some excess volatility). Next, we discuss a simple extension of the model where the exchange rate may disconnect from its fundamental value.

5.3 The Perception of the Fundamental Exchange Rate

This section explores how the agents perceive the fundamental exchange rate. Psychological evidence indicates that expectations are heavily influenced by the anchor and adjustment heuristic. Tversky and Kahneman (1974) report that in many situations, people make estimates by starting from an initial value that is adjusted to yield the final answer. The initial value, or starting point, may be suggested by the formulation of the problem, or it may be the result of a partial computation. In either case, adjustments are typically insufficient implying biased estimates toward initial values.

Two questions arise. What is the relevant anchor of the agents and how do they carry out the adjustment? It seems natural that the anchor depends on the previously perceived fundamental exchange rate. In addition, there exists evidence that the exchange rate itself may be part of the anchor. We call this self-confirmation. Self-confirmation is included because of the way technical analysis is conducted. In the

[7] This is also reflected in the famous GARCH models (Bollerslev 1986). Using a GARCH (1,1) specification, the ARCH term is estimated as 0.058 and the GARCH term as 0.908. Hence, the conditional variance depends slightly on most recent shocks (ARCH effect) and strongly on the temporary volatility (GARCH effect). Since the sum of both terms is nearly one, the volatility shocks are quite persistent.

words of Murphy (1999), a technician believes that anything that can possibly affect the price – fundamentally, politically, psychologically, or otherwise – is actually reflected in the price of that market. Therefore, a technician indirectly studies fundamentals: if prices are rising, fundamentals must be bullish. A study of price action is all that is required. According to Murphy, agents believe that the exchange rate itself reflects relevant information.

The adjustment procedure consists of two steps. First, agents react, of course, to the arrival of new information. However, the exact meaning of new information is not clear. Experiments even indicate that agents tend to systematically misperceive news. Due to the conservatism heuristic, individuals are slow to change their beliefs in the face of new evidence. Individuals update their posteriors in the right direction, but by too little in magnitude (Edwards 1968). This may lead to an underreaction to news. On the contrary, the representativeness heuristic highlights the phenomenon that people think they see patterns in truly random sequences. This may cause an overreaction to news. For instance, after a consistent series of good news for an asset, investors conclude that the past history is representative of an underlying growth potential. While a consistent pattern of high growth may be nothing more than a random draw for the asset, investors see order among chaos and infer from the in-sample growth path that the asset just keeps growing (Tversky and Kahneman 1974).

The second adjustment step covers experience-based feedback learning. Although such a procedure may yield a partial error correction, the adjustment is typically quite slow in time and small in magnitude. Psychologists give various reasons for the incomplete changes such as overestimation and overconfidence of the agents, avoidance of cognitive dissonance, or simply a lack of cognitive skills (Kahneman, Slovic and Tversky 1986).

To summarize, the agents use a mixture of the previously perceived fundamental exchange rate and the exchange rate itself as an anchor. The adjustment takes part in two steps, first, by an update according to new information, and second, by an error correction in form of feedback learning. The remainder of this section discusses the implications of the anchor and adjustment heuristic for the exchange rate dynamics. Afterwards, a more detailed analysis of what influences the distortion is given. To clarify the forces at work, the discussion of the perception process is divided into two sections: without and with learning adjustment.

5.3.1 Perception without Learning

Our first version of the perception of the fundamental exchange rate is

$$LogS_t^{FP} = \varepsilon LogS_{t-1}^{FP} + (1-\varepsilon)LogS_{t-1} + \kappa_t \eta_t, \tag{4a}$$

where ε indicates the impact of the previously perceived fundamental exchange rate and the exchange rate for the anchor. We assume a self-confirmation of 10 percent (ε=0.9). The misperception of the news is expressed through κ, where κ~U(0.25, 1.75). Clearly, on average the news is perceived correctly, but mistakes occur every period.[8]

Figure 4 shows a typical simulation run for the modified solution. Now, the exchange rate (solid line) fluctuates around its perceived fundamental value (dotted line). Because of self-confirmation, this value slightly changes in every period. Moreover, the exchange rate has the potential to move away from its equilibrium value (dashed line). The selection outcome for the trading rules is not affected through (4a).

This is a main difference to the work of Frankel and Froot (1986). Their dynamics are as follows. Caused by an initial shock, the exchange rate starts to shift away from its fundamental. Via self-fulfilling expectations, chartists are gaining prominence so that the bubble path is supported. However, when almost all agents are chartists the bubble dynamics automatically die out. Afterwards, there occurs a fundamentalists revival. The weight of the fundamentalists increases as the exchange rate converges towards its long-term equilibrium. With this approach, Frankel and Froot explain the bubble path of the US dollar in the eighties. Note that their weighting scheme changes slowly and smoothly during a (lasting) bubble, whereas in our setup the importance of the trading rules varies more rapidly and irregularly. In general, we find our scenario more realistic.

---

[8] Since E[κ]=1, news is not systematically misinterpreted. For this, compare section 5.3.3.

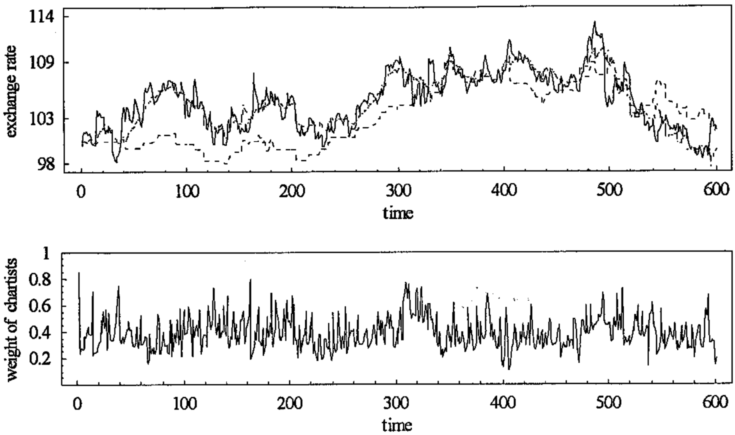

Figure 4: Exchange Rate Dynamics and Weight of Chartists (without Learning). The solid line is the exchange rate, the dashed line its fundamental, and the dotted line the perceived value of the fundamental. Parameters are as in table 1, additional $\varepsilon = 0.9$, $\kappa \sim U\ (0.25, 1.75)$.

Figure 5 displays a simulation run for the exchange rate and its fundamental over 40,000 periods, where the data is plotted every 80 periods. If one assumes that the model is based on daily observations, the entries (roughly) represent quarterly data. The model produces both periods where the exchange rate is close to its fundamental (around quarter 150) and periods of large bubbles (between quarter 280 and 320). In the latter case, the exchange rate is more than 30 percent overvalued. The patterns of the distortions can be quite different. For instance, around quarter 260 one finds a typical bubble. During a series of good news, the exchange rate overshoots its fundamental. Between quarter 90 and 110, the exchange rate is relatively stable whereas the fundamental shifts away. Finally, at quarter 230 the exchange rate and its fundamental even move in opposite directions. In the bottom part of figure 5, the current account reaction is given. Naturally, bubbles give rise to current account imbalances.

Figure 6 is designed to provide a first answer to the question of what may trigger a bubble. The top right shows a simulation run where firms′ transactions and the self-confirmation are both zero. Due to the misperception, the exchange rate starts to diverge from its fundamental without any long-term mean-reversion. In the bottom right, the self-confirmation is included again. Self-confirmation has the power to destabilize the

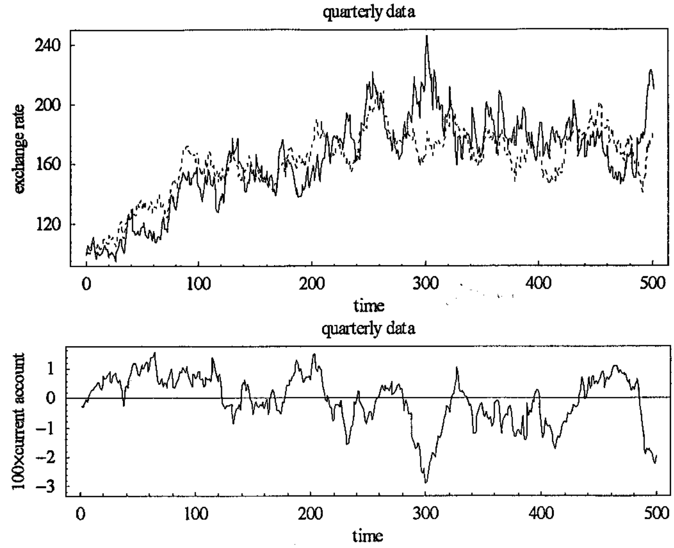

Figure 5: Exchange Rate Dynamics and Current Account Reaction (without Learning). The solid line is the exchange rate, the dashed line its fundamental. Parameters are as in table 1, additional $\varepsilon=0.9$, $\kappa \sim$ U (0.25, 1.75), 40,000 observations, plotted every 80 periods.

market even further. In the bottom left, the firms′ transactions are included but no self-confirmation. The demand of the firms is 40 times higher than before. In contrast to the top right, the exchange rate is pushed closer to its fundamental but still does not track it. The top left shows the base run again. Including both self-confirmation and trade transactions leads to a closer relationship between the exchange rate and its fundamental. The reason is that the exchange rate truly contains relevant information for the perception of the fundamental value. Self-confirmation behavior is not totally irrational. Whenever the difference between the exchange rate and its fundamental is small, self-confirmation increases the distortion. For example, misperceptions of news or feedback trading of the speculators may push the exchange rate away from its fundamental. These misalignments are settled via self-confirmation. However, with increasing misalignments, current account imbalances also rise. Period for period trade transactions induce some pressure on the exchange rate into the direction of its fundamental. Now, the bubble path is brought to an end via self-confirmation.

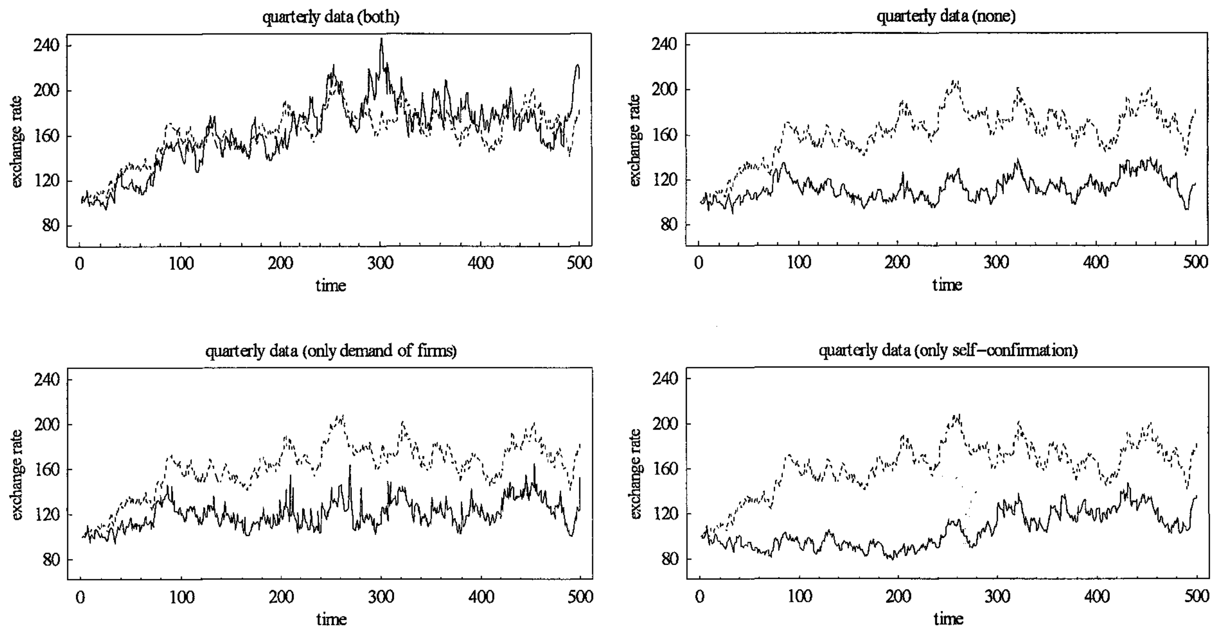

Figure 6: Comparison of Exchange Rate Dynamics (without Learning). The solid line is the exchange rate, the dashed line its fundamental. Top left contains a simulation with both firms' transactions and self-confirmation, top right without both of them, bottom left only demand from the firms (40 times higher than before), and bottom right only self-confirmation. Parameters are as in table 1, additional $\varepsilon = 0.9$, $\kappa \sim U(0.25, 1.75)$, 40,000 observations, plotted every 80 periods.

5.3.2 Perception with Learning

So far, the agents adjust their anchor with respect to the arrival of new information. However, the dynamics are able to generate huge current account imbalances. It seems natural that the agents try to learn from these macroeconomic disequilibria. The final version of the perception of the fundamental exchange rate is

$$LogS_t^{FP} = \varepsilon LogS_{t-1}^{FP} + (1-\varepsilon)LogS_{t-1} + \kappa_t \eta_t + Log[1+\omega d_{t-1}^{IT}], \quad (4b)$$

where ω indicates the degree of feedback learning. Note that the agents do not infer the true fundamental exchange rate out of the current account data, but simply adjust their anchor to some extent into the right direction.

Figure 7 and 8 display the short-run and long-run dynamics of the final system. The simulated time series resemble the previous ones. The exchange rate circles around its perceived equilibrium value and the weights of chartists fluctuate up and down in a range from 15 to 80 percent. Furthermore, the exchange rate has the power to drift away from its equilibrium value for some time, but eventually a mean-reversion reaction sets in. As a consequence of the feedback learning, the distortion is smaller.

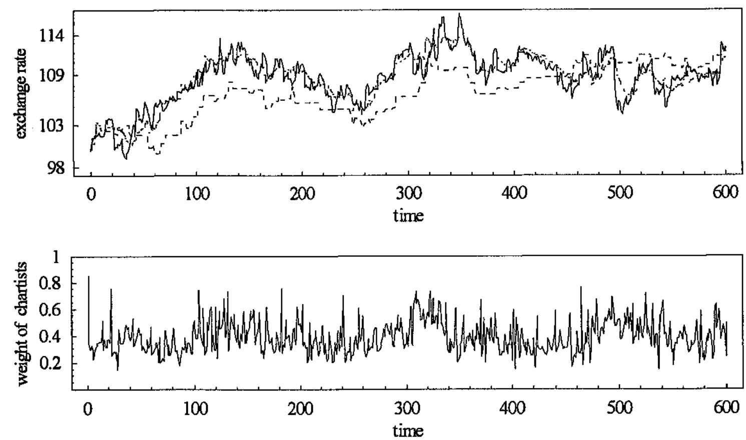

Figure 7: Exchange Rate Dynamics and Weight of Chartists (with Learning). The solid line is the exchange rate, the dashed line its fundamental, and the dotted line the perceived value of the fundamental. Parameters are as in table 1, additional $\varepsilon = 0.9$, $\kappa \sim U(0.25, 1.75)$, $\omega = 0.25$.

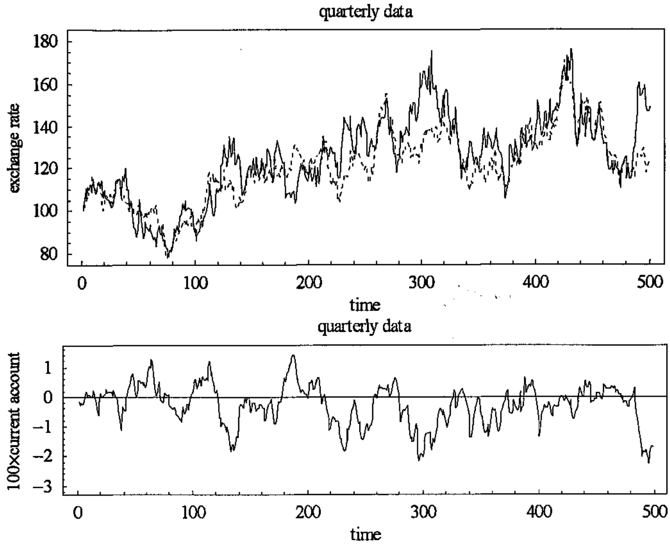

Figure 8: Exchange Rate Dynamics and Current Account Reaction (with Learning). The solid line is the exchange rate, the dashed line its fundamental. Parameters are as in table 1, additional $\varepsilon = 0.9$, $\kappa \sim U(0.25, 1.75)$, $\omega = 0.25$, 40,000 observations, plotted every 80 periods.

To many economists, the reliance on technical trading rules appears to be irrational. However, Huang and Day (1993) state that one must take note of the fact that during bear and bull markets, chartists are correct except at turning points. In this sense, they are right more often then they are wrong. This may be one of the reasons for the existence of chartists even if on average they lose wealth in the long-run. Next, we try to test this presumption with our data set. Suppose a small individual investor enters the market in period t=a and uses only the technical trading rule as specified by equation (1). Then, after T periods his profits are

$$P_T = S_T \sum_{t=a}^{T} d_t^C - \sum_{t=a}^{T} S_{t-1} d_t^C, \tag{11}$$

where the first term stands for the final revenue from clearing the position and the second term for the permanent revenue from building up the position. Figure 9 shows the evolution of the profits for the case in which the agent enters the market in period a=1 (top) and a=20,001 (bottom). In the long-run, technical analysis does not seem to be profitable. Nevertheless, one also finds longer periods where the agent earns money by simply applying his rule. For instance, if the agent starts to trade in period 20,001, then his profits are positive even after 20,000 periods. One reason for this phenomenon is that technical trading rules are especially successful in bubble times. In these periods, they tend to build up a larger position which increases in value as the bubble moves on (compare the last 1,000 periods). Thus, we stress that the use of simple technical trading rules as an adaptive scheme of behavior need not be totally irrational per se. Moreover, a number of empirical studies impressively demonstrates the success of a broad range of technical trading rules (Schulmeister 1988, Brock et al. 1992).

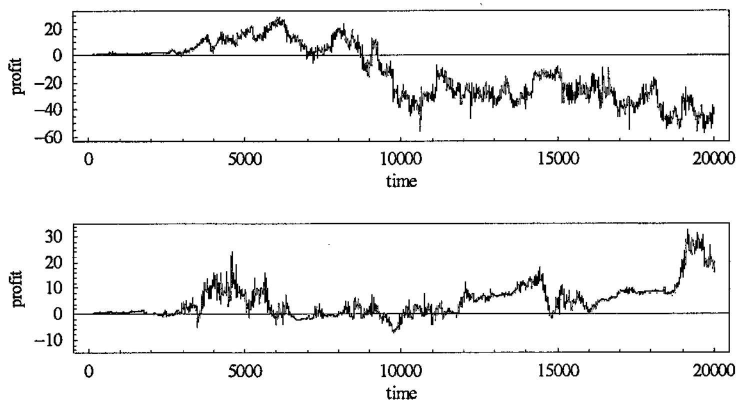

Figure 9: The Profitability of Technical Trading Rules. Parameters are as in table 1, additional ε=0.9, $\kappa \sim$ U(0.25, 1.75), ω= 0.25, in the top profits from periods 1-20,000, in the bottom from 20,001-40,000.

In addition, figure 8 seems to indicate that the foreign exchange market is rather inefficient. A cointegration analysis allows us to check this presumption. Variables are cointegrated if there exists a linear combination that is stationary (Engle and Granger 1987). If the exchange rate tracks its fundamental, then the difference between the exchange rate and its fundamental should be stationary. In a broader sense, an equilibrium exists if the difference between the variables does not become too large.

Cointegration may be tested according to the following procedure (Enders 1994). The first step consists of testing the order of integration. By definition, cointegration necessitates that the variables to be integrated are of the same order. For this pre-test, the augmented Dickey-Fuller (ADF) test can be applied. If both variables are integrated of the order 1, the next step is to estimate the long-run equilibrium relationship between S^F and S.[9] For this, one has to regress

$$S_t^F = aS_t + u_t, \tag{12}$$

and then apply the ADF test on the residuals u. Using a four-lag specification delivers

$$\Delta u_t = a_0 u_{t-1} + a_1 \Delta u_{t-1} + a_2 \Delta u_{t-2} + a_3 \Delta u_{t-3} + a_4 \Delta u_{t-4}. \tag{13}$$

If it is not possible to reject the null hypothesis $a_0=0$, the hypothesis that the variables are not cointegrated cannot be rejected. In simpler terms, if S and S^F were found to be integrated of order 1 and the residuals are stationary, one can conclude that the series are cointegrated.

Table 3 reports the cointegration results for the following simulation design. Subsamples are obtained by dividing a total number of 40,000 observations by a maximal time period. For instance, using a maximum time period of 1,000 one obtains 40 subsamples. For each of them, a unit root and cointegration test is carried out with a significance level of 5 percent. Table 3 indicates that the price tracks its fundamental value only in the long-run. For shorter time periods, there seems to be no cointegration. For instance, if the maximum number of observations is 500, only 35 percent of the subsamples show an equilibrium relationship between the exchange rate and its fundamental value. In other words, in 65 percent of the cases, the exchange rate moves away from its long-term equilibrium, which should be interpreted as a clear sign of market inefficiency.[10]

Table 4 highlights a problem empirical cointegration studies face. A total of 100,000 observations is separated into 10 subsamples. Cointegration results are presented for different degrees of time aggregation. For instance, a time aggregation of 5 means that the variable under consideration is available only every 5 trading periods. Thus, for each subsample only 2000 observations remain. Table 4 demonstrates that it becomes increasingly difficult to spot cointegration relationships if the time aggregation is too large. While the test delivers 9 significant cointegration relationships for a time

[9] Variables are integrated of the order 1 if the first differences of the variables are stationary. The level of the exchange rates is, in general, not stationary (unit roots), else the integration order would be 0.

[10] Although almost all subsamples display unit roots in the exchange rates, the empirical t-values of the ADF test increase if the sample length decreases. For time horizons below 100 periods, the exchange rates sometimes appear to be integrated of order 0.

aggregation of 1, these fall to zero if the time aggregation is 20. And this, although the fundamental exchange rate is known exactly and does not have to be calculated out of an (biased) equilibrium model.[11]

time horizon	40,000	20,000	10,000	5,000	4,000	2,000	1,000	500	250
number of subsamples	1	2	4	8	10	20	40	80	160
number of unit roots in S	1	2	4	8	10	20	40	80	159
number of cointegration	1	2	4	5	5	8	13	28	39
percentage of cointegration	100	100	100	62.5	50	40	32.5	35	24.4

Table 3: Cointegration Analysis between the Exchange Rate and its Fundamental. Parameters are as in table 1, additional $\varepsilon = 0.9$, $\kappa \sim U\ (0.25, 1.75)$, $\omega = 0.25$, 40,000 observations. The significance level is 5 percent.

subsambles	10	10	10	10	10	10
time aggregation	1	2	4	5	10	20
observations	10,000	5,000	2,500	2,000	1,000	500
number of unit roots in S	10	10	10	10	10	10
number of cointegration	9	5	2	1	1	0

Table 4: Cointegration Analysis under Time Aggregation. Parameters are as in table 1, additional ε=0.9, $\kappa \sim$ U(0.25, 1.75), $\omega = 0.25$, 100,000 observations. The significance level is 5 percent.

The main characteristic time series properties are, in general, not affected by the modifications. Table 5 summarizes some descriptive statistics of the returns for a simulation run over 10,000 periods (compare with table 2). In particular, the most extreme exchange rate movements and the variance of the returns behave well. The kurtosis is clearly higher than 3.

[11] In the empirical literature, evidence for cointegration between the exchange rate and its long-run equilibrium (derived from structural relationships such as a monetary model) is weak (Taylor 1995).

min	median	max	variance	skewness	kurtosis
7.56 %	0.00 %	7.56 %	0.000068	0.052	11.01

Table 5: Descriptive Statistics of Returns (with Learning). Parameters are as in table 1, additional ε=0.9, $\kappa \sim$ U (0.25, 1.75), ω = 0.25, 10,000 observations.

Furthermore, the overall shape of the distribution of the returns, as visible from figure 10, matches the empirical one (compare with figure 2). For example, taking the largest 30 percent of the observations into account, the tail index is estimated as 3.54. Again, this indicates fat tails.

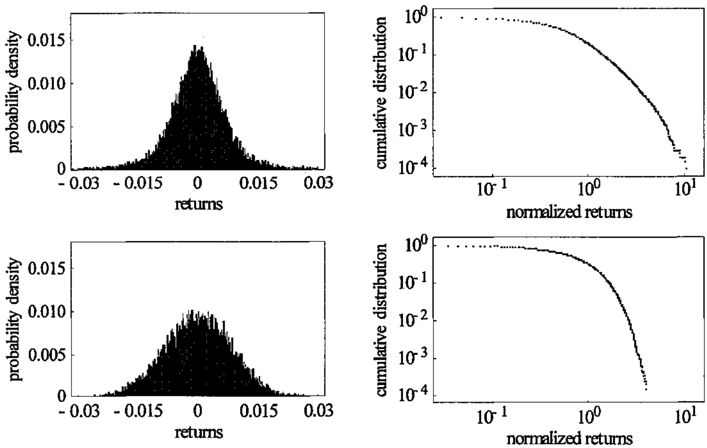

Figure 10: Distribution of Returns and Scaling Behavior (with Learning). The top shows the distribution of the returns and the scaling behavior of the cumulative distribution of the positive and negative tails for normalized Log-returns (parameters as in table 1, additional ε = 0.9, $\kappa \sim$ U(0.25, 1.75), ω = 0.25, 20,000 observations), the bottom the same but for a Normal distribution with identical variance.

Figure 11 displays the evolution of the returns and the trading volume together with some autocorrelation functions (compare with figure 3). One finds almost no autocorrelation for raw returns, but persistent volatility clustering for absolute returns and trading volume. Finally, the cross-correlation between volatility and total transactions is high for current volumes and low for past volumes.

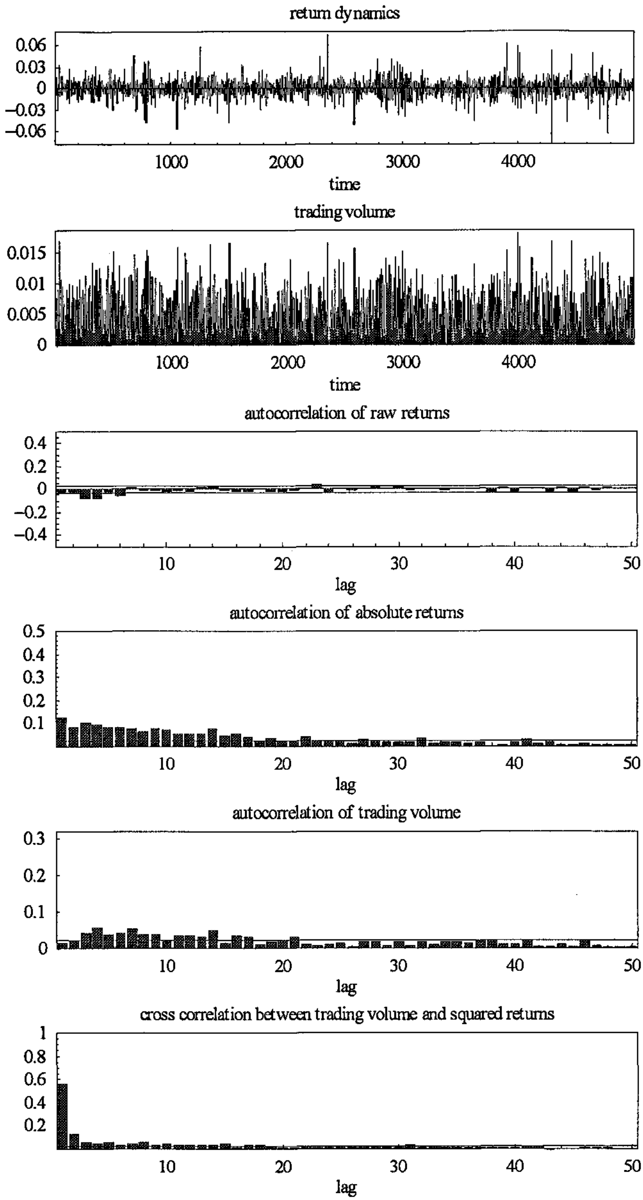

Figure 11: Volatility Clustering (with Learning). Parameters are as in table 1, additional $\varepsilon = 0.9$, $\kappa \sim U(0.25, 1.75)$, $\omega = 0.25$, T=5,000 observations. 95 percent confidence intervals are plotted as $\pm 2/\sqrt{T}$ (assumption of white noise).

5.3.3 Analysis of Distortion

In our model, the distortions are caused by the way the speculators perceive the fundamental exchange rate. This phenomenon is robust in the sense that it is qualitatively not affected by different values of the parameters. However, by varying a single coefficient, quantitative changes emerge. Focussing on these changes, this section tries to develop a deeper and more systematic understanding of the driving forces of the bubble processes. To be more accurate, a distortion measure is defined as

$$D = \frac{100}{T}\sum_{t=1}^{T}\left|(S_t - S_t^F)/S_t^F\right| , \tag{14}$$

where the distortion D is computed as the average relative distance between the exchange rate and its fundamental.

Figure 12 presents some results for the following simulation exercise. The distortion is calculated out of 40,000 data points for each entry in one part of the figure. The time series are computed with the same parameter setting as before. Moreover, the same seed for the random variables is used. The parameter under consideration is varied as marked on the axis.

What are the results? A high systematic reaction of the technical trading rule $\alpha^{C,1}$ goes along with an increase in the distortion. The more systematically the technicians trade, the more often they induce a temporary trend into the time series. Shocks, such as a strong misperception of new information or a high random demand, yield an exchange rate movement which is amplified and prolonged into the next periods.

The expected adjustment speed of the exchange rate towards its perceived fundamental γ is randomly drawn from a Uniform distribution within the bounds of 0 to 0.5. In the top right of figure 11, γ is multiplied by a factor ranging from 1 to 2. One sees that the higher the adjustment speed, the lower the distortion. The explanation is that for a low γ, the expected future exchange rate is not influenced strongly through its fundamental, so that the mean-reversion is weak. Moreover, if the probability of random changes of γ increases, the distortion shrinks further (not displayed). Lasting sequences of low and high γ are more destabilizing than a stronger mix of γ.

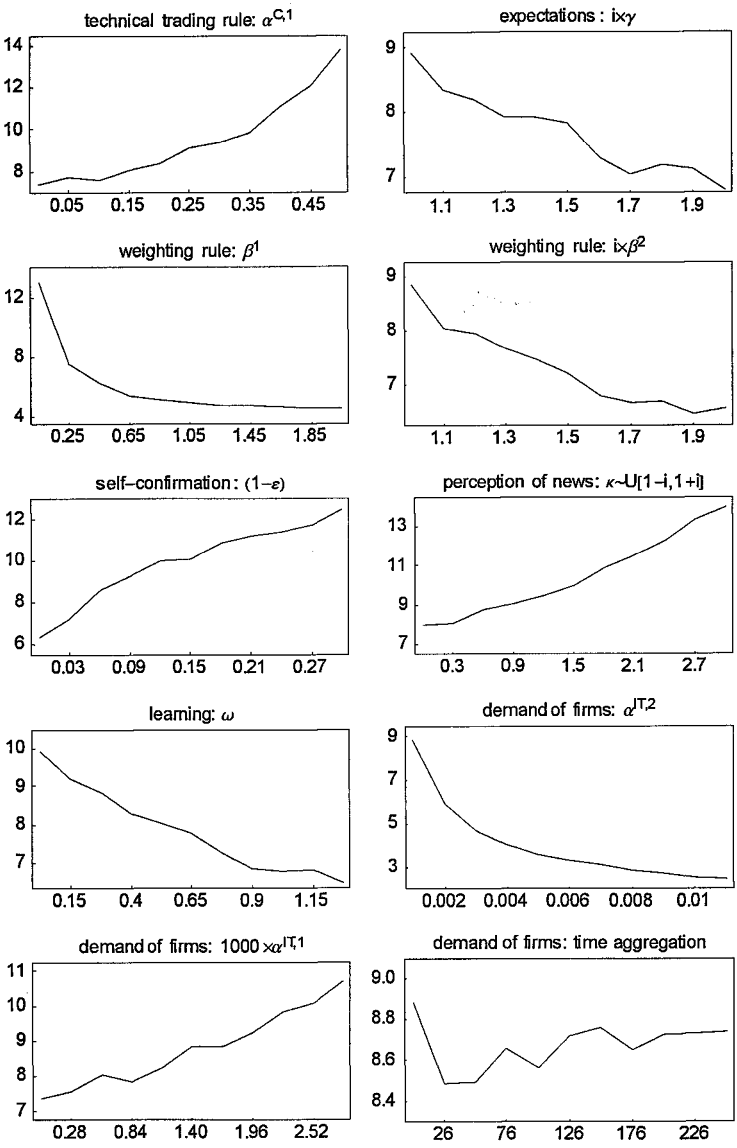

Figure 12: Analysis of Distortion. Parameters are as in table 1, additional $\varepsilon = 0.9$, $\kappa \sim U(0.25, 1.75)$, $\omega = 0.25$, T = 40,000 observations. Parameters are varied as marked on the axis.

The next two parts of figure 12 demonstrate that an increase in the basic fraction of fundamentalists β^1 or an increase in the popularity of the fundamental trading rule β^2 reduces the distortion. The higher the degree of fundamentalism, the lower the exchange rate movements. This reduces the demand from technical traders, since their trading signals are less pronounced. As a consequence, their behavior becomes less trend oriented. In contrast to the technical traders, the fundamentalists stabilize the foreign exchange market.

As discussed in section 5.3.1 (figure 6), the implications of self-confirmation for the distortion appear to be ambiguous (without learning). On the one hand, including the spot rate into the anchor destabilizes the market since misalignments are settled into the course. On the other hand, if the current account is not balanced, the exchange rate may transport relevant information so that self-confirmation can stop bubbles. The overall effect of self-confirmation is revealed in figure 12. The distortion is positively correlated with self-confirmation.

The case of misperception is straightforward. The higher the misperception, the higher the distortion. Naturally, the distortion increases further, if the perception of the news is biased (not displayed). Periods of overreaction to news lead to traditional bubbles (overshooting), whereas in periods of underreaction to news the exchange rate does not follow the development of its fundamental.[12]

A stronger (error correction) learning on the part of the agents yields a lower distortion. However, the effectiveness of the learning behavior also depends on the clarity of the signal they have. A higher current account reaction to misalignments $\alpha^{IT,2}$ diminishes the distortion. First, it induces a stronger permanent pressure on the exchange rate. Second, the exchange rate is driven back to its fundamental via self-confirmation. Third, the higher $\alpha^{IT,2}$, the more often the agents receive the true learning signal. The opposite occurs if the random component of the firms′ demand $\alpha^{IT,1}$ rises.

[12] See Barberis, Shleifer and Vishny (1998) for an analytical exposition on systematic misperception.

Finally, in our model the agents use daily trade transactions of the firms as the learning signal. The last part of the figure shows the result when the signal is given as an aggregate over some time. A time aggregation of 76 means that the agents only have access to firms′ demand quarterly (the average over 76 trading days). Surprisingly, inferring macroeconomic imbalances from different time horizons does not seem to be that relevant. Although time aggregation makes the signal less noisy (since the random components cancel out), the signal arrives with a time lag. One observes a slightly higher distortion only for very long periods, say over 300 trading periods.

5.4 Conclusions

To sum up, the aim of this chapter is threefold: to develop a model that gives rise to bubbles but also delivers realistic exchange rate movements, to explore how the agents perceive the fundamental exchange rate, and to unravel the forces behind the distortions. The trading activity of the agents is described with the aid of an empirical micro-foundation. The interaction between the traders generates complex dynamics which display unit roots in the exchange rates, fat tails for returns and volatility clustering.

Fundamental to this chapter is the perception of the fundamental exchange rate. Psychologists claim that the agents behave according to the anchor and adjustment heuristic. Since the adjustment is typically incomplete and mistakes are propagated over time, the perceived value of the fundamental can strongly deviate from its true value even for longer periods. The degree of distortions depends on several forces. For instance, it increases with a higher systematic behavior of the chartists, or decreases with a higher popularity of the fundamentalists.

This chapter is built on empirical evidence. However, we understand this chapter only as one step to achieving a better linking between economic modeling and psychological evidence. What is needed is a better understanding of the human cognition process. How do the agents perceive and process information, how do they select an action, and how do they learn. Recent experiments along the lines of Hommes et al. (1999) and Sonnemans et al. (1999) hopefully provide better and more detailed empirical grounds for future research.

6 Concluding Remarks

What drives exchange rates? This study has identified several factors: the trading activity of the speculators, the news arrival process, transactions of international firms, and central bank interventions.

- The nonlinear interaction between technical and fundamental trading rules is seen as a strong endogenous force of exchange rate variability.
- Exchange rates react to the arrival of news in two ways. On the one hand, the fundamental value of the exchange rate changes. On the other hand, news is often misperceived. Even if the perception process is not systematically biased, lasting deviations of the exchange rate from its fundamental value are likely.
- In the short run, transactions of international firms appear to be random and enable exchange rates to tremble. In the long run, the current account adjustment tends to bring the exchange rate back to its fundamentals.
- Central bank interventions also influence the time series properties of exchange rates. For example, periods of low and high volatility may occur if the central bank alternates between different kinds of intervention operations.

It is heavily debated whether the foreign exchange market is efficient or not. We have found evidence that exchange rates are excessively volatile and may diverge from fundamental values. A higher exchange rate volatility obviously implies a higher risk for the firms engaged in international trade. Since exchange rate risk can only partially be managed, international trade is hampered and thus physical investment depressed. Due to lasting distortions new investments may flow into the wrong projects. If the market corrects this misalignment, the profitability of the investment may vanish. From this point of view, the foreign exchange market seems to be inefficient.

However, our analysis suggests two channels which may reduce market inefficiencies. Since fundamentalists tend to stabilize the dynamics, an increase in the popularity of fundamentalism may reduce the volatility. This may be realized by providing better information about the future state of the economy. The effectiveness of central bank interventions depends on market circumstances. If the demand of the chartists is correlated, a "leaning against the wind" operation may calm down the market. If the traders are uncertain about fundamentals, the central bank may reduce distortions by supporting the fundamental exchange rate. But it should not be overlooked that the conditions for successful intervention are often not fulfilled. Then central bank intervention can increase market instability. Unfortunately, alternative means to stabilize the foreign exchange market, such as a transaction tax or price controls, have not yet been thoroughly explored.

This leads us to the question about the future direction of the chartists-fundamentalists approach. In the last decade, several mechanisms were proposed to model exchange rate behavior. Some of them favor the view that dynamics are generated by the nonlinear nature of the trading rules. Others see the complex switching process between different kinds of trading rules due to profit differentials or social interactions as the driving force of exchange rate fluctuations. By now, this area of research has made considerable progress; the basic principles of the dynamics are quite well understood.

Contemporary research starts to model the behavior of the agents in more detail. Different degrees of memory and learning capabilities as well as psychological aspects like emotions are increasingly investigated. Other extensions try to implement a chartists-fundamentalists setup into a larger macroeconomic framework to endogenize the evolution of the fundamental exchange rate. By this, one aims to study how the variability of exchange rates is transmitted to other markets and how this influences macroeconomic variables like income or employment.

Besides a positive analysis of foreign exchange dynamics, the first normative issues are raised: chartists-fundamentalists models can be used as a laboratory to compare various means of market stabilization. Focussing on the market microstructure, different organizational aspects of speculative markets may be evaluated. Our contribution towards central bank intervention can be seen as a preliminary step in this direction. However, this strand of literature is newborn. The future will show if this kind of research is able to improve the functioning of the markets.

References

Almeida, A., Goodhart, C. and Payne, R. (1998): The Effects of Macroeconomic News on High Frequency Exchange Rate Behavior, Journal of Financial and Quantitative Analysis, 33, 383-408.

Andersen, T. and Bollerslev, T. (1997): Intraday Periodicity and Volatility Persistence in Financial Markets, Journal of Empirical Finance, 4, 115-158.

Andersen, T. and Bollerslev, T. (1998): Deutsche Mark-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer Run Dependencies, Journal of Finance, 53, 219-265.

Arifovic, J. (1996): The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies, Journal of Political Economy, 104, 510-541.

Arthur, B., Holland, J., LeBaron, B., Palmer, R. and Tayler, P. (1997): Asset Pricing under Endogenous Expectations in an Artificial Stock Market, in: Arthur, B., Durlauf, S. and Lane, D. (Eds.): The Economy as an Evolving Complex System II, Addison-Wesley: Reading, 15-44.

Bajo-Rubio, O., Fernandez-Rodriguez, F. and Sosvilla-Rivero, S. (1995): Chaotic Behavior in Exchange Rate Series: First Results for the Peseta-US Dollar Case, in: Trippi, R. (Ed.): Chaos and Nonlinear Dynamics in the Financial Markets, Irwin Concorde: Ontario, 327-334.

Bajoumi, T. (2001): The Morning After: Explaining the Slowdown in Japanese Growth in the 1990s, Journal of International Economics, 53, 241-259.

Banerjee, A. (1992): A Simple Model of Herd Behavior, Quarterly Journal of Economics, 107, 797-817.

Barberis, N., Shleifer, A. and Vishny, R. (1998): A Model of Investor Sentiment, Journal of Financial Economics, 49, 307-343.

Barnett, W. and Serletis, A. (2000): Martingales, Nonlinearity, and Chaos, Journal of Economic Dynamics and Control, 24, 703-724.

Bartolini, L. and Giorgianni, L. (2000): Excess Volatility of Exchange Rates with Unobservable Fundamentals, Federal Reserve Bank of New York, Staff Reports Number 103.

Baumol, W. (1957): Speculation, Profitability, and Stability, Review of Economics and Statistics, 39, 263-271.

Baumol, W. and Benhabib, J. (1989): Chaos: Significance, Mechanism, and Economic Applications, Journal of Economic Perspectives, 3, 77-105.

Beja, A. and Goldman, B. (1980): On the Dynamic Behavior of Prices in Disequilibrium, Journal of Finance, 35, 235-248.

BIS (1999): Central Bank Survey of Foreign Exchange and Derivatives Market Activity 1998, Basle.

Black, F. (1986): Noise, Journal of Finance, 41, 529-543.

Bollerslev, T. (1986): Generalized Autoregressive Conditional Heteroskedasticity, Journal of Econometrics, 31, 307-327.

Brock, W., Lakonishok, J. and LeBaron, B. (1992): Simple Technical Trading Rules and the Stochastic Properties of Stock Returns, Journal of Finance, 67, 1731-1764.

Brock, W. and Potter, S. (1993): Nonlinear Time Series and Macroeconometrics, in: Maddala, G., Rao, C. and Vinod, H. (Eds.): Handbook of Statistics, 11, North-Holland: Amsterdam,195-229.

Brock, W. and LeBaron, B. (1996): A Dynamic Structural Model for Stock Return Volatility and Trading Volume, Review of Economics and Statistics, 78, 94-110.
Brock, W. and Hommes, C. (1997a): A Rational Route to Randomness, Econometrica, 65, 1059-1095.
Brock, W. and Hommes, C. (1997b): Models of Complexity in Economics and Finance, in: Heij, C., Schumacher, J., Hanzon, B. and Praagman, C. (Eds.): System Dynamics in Economic and Financial Models, Wiley: New York, 3-41.
Brock, W. and Hommes, C. (1998): Heterogeneous Beliefs and Routes to Chaos in a Simple Asset Pricing Model, Journal Of Economic Dynamics Control, 22, 1235-1274.
Brock, W. and Hommes, C. (1999): Rational Animal Spirits, in: Herings, P., Talman, A. and Laan, G. van der (Eds.): The Theory of Markets, North-Holland: Amsterdam, 109-137.
Cai, J., Cheung, Y., Lee, R. and Melvin, M. (2001): "Once-in-a-Generation" Yen Volatility in 1998: Fundamentals, Intervention, and Order Flow, Journal of International Money and Finance, 20, 327-347.
Camerer, C. (1995): Individual Decision Making, in: Kagel, J. and Roth, A. (Eds.): The Handbook of Experimental Economics, Princeton University Press: Princeton, 587-703.
Campbell, J., Lo, A. and MacKinlay, C. (1997): The Econometrics of Financial Markets, Princeton University Press: Princeton.
Cheung, Y. and Wong, C. (2000): A Survey of Market Practitioners's Views on Exchange Rate Dynamics, Journal of International Economics, 51, 401-419.
Cheung, Y. and Chinn, M. (2001): Currency Traders and Exchange Rate Dynamics: A Survey of the US Market, Journal of International Money and Finance (in press).
Chiarella, C. (1992): The Dynamics of Speculative Behavior, Annals of Operations Research, 37, 101-123.
Conlisk, J. (1996): Why Bounded Rationality? Journal of Economic Literature, 34, 669-700.
Curcio, R., Goodhart, C., Guillaume, D. and Payne, R. (1997): Do Technical Trading Rules Generate Profits? Conclusions from the Intra-Day Foreign Exchange Market, International Journal of Finance and Economics, 2, 267-280.
Cutler, D., Poterba, J. and Summers, L. (1990): Speculative Dynamics and the Role of Feedback Traders, American Economic Review, 80, 63-68.
Cutler, D., Poterba, J. and Summers, L. (1991): Speculative Dynamics, Review of Economic Studies, 58, 529-546.
Da Silva, S. (1999): Exchange Rate Dynamics Redux and Chaos, Birmingham University Department of Economics Discussion Papers, No. 99-08
Day, R. and Huang, W. (1990): Bulls, Bears and Market Sheep, Journal of Economic Behavior and Organization, 14, 299-329.
De Bondt, W. (1998): A Portrait of the Individual Investor, European Economic Review, 42, 831-844.
De Grauwe, P. and Dewachter, H. (1992): Chaos in the Dornbush Model of the Exchange Rate, Kredit und Kapital, 25, 26-54.
De Grauwe, P. and Dewachter, H. (1993): A Chaotic Model of the Exchange Rate: The Role of Fundamentalists and Chartists, Open Economies Review, 4, 351-379.
De Grauwe, P., Dewachter, H. and Embrechts, M. (1993): Exchange Rate Theory - Chaotic Models of Foreign Exchange Markets, Blackwell: Oxford.

De Long, B., Shleifer, A., Summers, L. and Waldmann, R. (1990a): Positive Feedback Investment Strategies and Destabilizing Rational Speculation, Journal of Finance, 45, 379-395.

De Long, B., Shleifer, A., Summers, L. and Waldmann, R. (1990b): Noise Trader Risk in Financial Markets, Journal of Political Economy, 98, 703-738.

Deutsche Bundesbank (1998): Devisenkursstatistik, Frankfurt.

Dickey, D. and Fuller, W. (1979): Distribution of the Estimators for Autoregressive Time Series with a Unit Root, Journal of the American Statistical Association, 74, 427-431.

Dickey, D. and Fuller, W. (1981): Likelihood Ratio Test Statistics for Autoregressive Time Series with a Unit Root, Econometrica, 49, 1057-1072.

Dominguez, K. and Frankel, J. (1993): Does Foreign Exchange Intervention Work? Institute of International Economics, Washington.

Dooley, M. and Shafer, J. (1983): Analysis of Short-Run Exchange Rate Behavior: March 1973 to November 1981, in: Bigman, D. and Taya, T. (Eds.): Exchange Rate and Trade Instability: Causes, Consequences and Remedies, Ballinger: Cambridge, 43-69.

Dunis, C. (1989): Computerized Technical Systems and Exchange Rate Movements, in: Dunis, C. and Feeny, M. (Eds): Exchange Rate Forecasting, Woodhead-Faulkner: New York, 165-205.

Edison, H. (1993): The Effectiveness of Central-Bank Intervention: A Survey of the Literature After 1982, Special Papers in International Economics, Number 18, International Finance Section, Princeton University.

Edwards, R. and Magee, J. (1966): Technical Analysis of Stock Trends, John Magee: Boston.

Edwards, W. (1968): Conservatism in Human Information Processing, in: Kleinmuntz, B. (Ed.): Formal Representation of Human Judgment, Wiley: New York, 17-52.

Eichenbaum, M. and Evans, C. (1995): Some Empirical Evidence on the Effects of Shocks to Monetary Policy on Exchange Rates, Quarterly Journal of Economics, 110, 975-1009.

Eichengreen, B., Tobin, J. and Wyplosz, C. (1995): Two Cases for Sand in the Wheels of International Finance, Economic Journal, 105, 162-172.

Enders, W. (1995): Applied Econometric Time Series, Wiley: New York.

Engle, R. (1982): Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation, Econometrica, 50, 987-1007.

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.

Farmer, D. (2000): Market Force, Ecology, and Evolution, Discussion Paper Santa Fe Institute, 98-12-117E.

Feeny, M. (1989): Charting the Foreign Exchange Markets, in: Dunis, C. and Feeny, M. (Eds): Exchange Rate Forecasting, Woodhead-Faulkner: New York, 99-163.

Fernandez-Rodriguez, F., Gonzales-Martel, C. and Sosvilla-Rivero, S. (2000): On the Profitability of Technical Trading Rules Based on Artificial Neural Networks: Evidence from the Madrid Stock Market, Economics Letters, 69, 89-94.

Flood, M. (1991): Microstructure Theory and the Foreign Exchange Market, Federal Reserve Bank of St. Louis Review, 73, 52-70.

Frank, M. and Stengos, T. (1989): Measuring the Strangeness of Gold and Silver Rates of Return, Review of Economic Studies, 56, 553-567.

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.

Frankel, J. and Froot, K. (1990): Chartists, Fundamentalists and the Demand for Dollars, in: Courakis, A. and Taylor, M. (Eds): Private Behavior and Government Policy in Interdependent Economies, Claendon Press: Oxford, 73-126.

Frankel, J. and Rose, A. (1995): Empirical Research on Nominal Exchange Rates, in: Grossman, G. and Rogoff, K. (Eds.): Handbook of International Economics, 3, Elsevier: Amsterdam, 1689-1729.

Frankel, J., Galli, G. and Giovannini, A. (1996): The Microstructure of Foreign Exchange Markets, University of Chicago Press: Chicago.

Frenkel, M. (1997): Exchange Rate Dynamics with Chartists, Fundamentalists, and Rational Speculators in the Foreign Exchange Market, International Journal of Business, 2, 1-24.

Friedman, M. (1953): The Case for Flexible Exchange Rates, in: Friedman, M. (Ed.): Essays in Positive Economics, University of Chicago Press: Chicago, 157-203.

Gaunersdorfer, A. and Hommes, C. (2000): A Nonlinear Structural Model for Volatility Clustering, University of Amsterdam, CeNDEF Working Paper 00-02.

Goodhart, C. (1988): The Foreign Exchange Market: A Random Walk with a Dragging Anchor, Economica, 55, 437-460.

Goodhart, C. and Figliuoli, L. (1991): Every Minute Counts in Financial Markets, Journal of International Money and Finance, 10, 23-52.

Goodhart, C. and Taylor, M. (1992): Why Don´t Individuals Speculate in the Forward Foreign Exchange Market, Scottish Journal of Political Economy, 39, 1-13.

Goodhart, C. and Giugale, M. (1993): From Hour to Hour in the Foreign Exchange Market, Manchester School, 61, 1-34.

Goodhart, C., McMahon, P. and Ngama, Y. (1993a): Testing for Unit Roots with Very High Frequency Spot Exchange Rate Data, Journal of Macroeconomics, 15, 423-438.

Goodhart, C., Hall, S., Henry, S. and Pesaran, B. (1993b): News Effects in a High-Frequency Model of the Sterling-Dollar Exchange Rate, Journal of Applied Econometrics, 8, 1-13.

Goodhart, C. and O´Hara, M. (1997): High Frequency Data in Financial Markets: Issues and Applications, Journal of Empirical Finance, 4, 73-114.

Grassberger, P. and Procaccia, I. (1984): Dimensions and Entropies of Strange Attractors from a Fluctuating Dynamic Approach, Physica D, 13, 34-54.

Groen, J. (1999): Long Horizon Predictability of Exchange Rates: Is It for Real? Empirical Economics, 24, 451-469.

Guillaume, D. (1995): A Low-Dimensional Fractal Attractor in the Foreign Exchange Markets? in: Trippi, R. (Ed.): Chaos and Nonlinear Dynamics in the Financial Markets, Chicago, Irwin Concorde: Ontario, 269-294.

Guillaume, D., Dacorogna, M. Dave, R., Müller, 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.

Hamilton, J. (1994): Time Series Analysis, Princeton University Press: Princeton.

Heiner, R. (1983): The Origin of Predictable Behavior, American Economic Review, 73, 560-595.

Hilborn, R. (2000): Chaos and Nonlinear Dynamics, Second Edition, Oxford University Press: Oxford.

Hill, B. (1975): A Simple General Approach to Inference About the Tail of a Distribution, Annals of Statistics, 3, 1163-1173.

Holyoak, K. and Nisbett, R. (1988): Induction, in: Sternberg, R. and Smith, E. (Eds.): The Psychology of Human Thought, Cambridge University Press: Cambridge, 50-91.
Hommes, C., Sonnemans, J., Tuinstra, J. and Velden, H. van de (1999): Expectation Driven Price Volatility in an Experimental Cobweb Economy, University of Amsterdam, CeNDEF Working Paper 99-07.
Hommes, C. (2001): Financial Markets as Nonlinear Adaptive Evolutionary Systems. Quantitative Finance, 1, 149-167.
Huang, W. and Day, R. (1993): Chaotically Switching Bear and Bull Markets: The Derivation of Stock Price Distributions from Behavioral Rules, in: Day, R. and Chen, P. (Eds.): Nonlinear Dynamics and Evolutionary Economics, Oxford University Press: Oxford, 169-182.
Hung, J. (1997): Intervention Strategies and Exchange Rate Volatility: A Noise Trading Perspective, Journal of International Money and Finance, 16, 779-793.
Hurst, H. (1951): Long Term Storage Capacity of Reservoirs, Transactions of the American Society of Civil Engineers, 116, 770-799.
Ito, T. (1990): Foreign Exchange Rate Expectations: Micro Survey Data, American Economic Review, 80, 434-449.
Kahneman, D., Slovic, P. and Tversky, A. (1986): Judgment under Uncertainty: Heuristics and Biases, Cambridge University Press: Cambridge.
Kantz, H. and Schreiber, T. (1999): Nonlinear Time Series Analysis, Cambridge University Press: Cambridge.
Keynes, J.M. (1936): The General Theory of Employment, Interest, and Money, Harcourt Brace: New York.
Kirman, A. (1991): Epidemics of Opinion and Speculative Bubbles in Financial Markets, in: Taylor, M. (Ed.): Money and Financial Markets, Blackwell: Oxford, 354-368.
Kirman, A. (1993): Ants, Rationality, and Recruitment, Quarterly Journal of Economics, 108, 137-156.
Lawrenz, C. and Westerhoff, F. (2001): Modeling Exchange Rate Behavior with a Genetic Algorithm, Computational Economics (in press).
LeBaron, B. (1999): Technical Trading Rule Profitability and Foreign Exchange Intervention, Journal of International Economics, 49, 125-143.
LeBaron, B., Arthur, B. and Palmer, R. (1999): Time Series Properties of an Artificial Stock Market, Journal of Economic Dynamics and Control, 23, 1487-1516.
LeBaron, B. (2000): Agent Based Computational Finance: Suggested Readings and Early Research, Journal of Economic Dynamics and Control, 24, 679-702.
Levich R. and Thomas, L. (1993): The Significance of Technical Trading Rules Profits in the Foreign Exchange Market: A Bootstrap Approach, Journal of International Money and Finance, 12, 451-474.
Lo, A., Mamaysky, H. and Wang, J. (2000): Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation, Journal of Finance, 55, 1705-1765.
Lui, Y.-H. and Mole, D. (1998): The Use of Fundamental and Technical Analysis by Foreign Exchange Dealers: Hong Kong Evidence, Journal of International Money and Finance, 17, 535-545.
Lux, T. (1995): Herd Behavior, Bubbles and Crashes, Economic Journal, 105, 881-896.
Lux, T. (1997): Time Variation of Second Moments from a Noise Trader/Infection Model, Journal of Economic Dynamics and Control, 22, 1-38.
Lux, T. (1998): The Socio-Economic Dynamics of Speculative Markets: Interacting Agents, Chaos, and the Fat Tails of Return Distributions, Journal of Economic Behavior and Organization, 33, 143-165.

Lux, T. and Ausloos, M. (2000): Market Fluctuations I: Scaling, Multi-Scaling and Their Possible Origins, in: Bunde, A. and Schellnhuber, H. (Eds.): Theories of Desaster: Scaling Laws Governing Weather, Body and Stock Market Dynamics, Springer: Berlin.

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

Lux, T. (2001): The Limiting Extremal Behavior of Speculative Returns: An Analysis of Intra-Daily Data from the Frankfurt Stock Exchange, Applied Financial Economics (in press).

Lyons, R. (2001): The Microstructure Approach to Exchange Rates, MIT Press: Cambridge.

MacDonald, R. (1999): Exchange Rate Behavior: Are Fundamentals Important? Economic Journal, 109, 673-691.

Mandelbrot, B. (1963): The Variation of Certain Speculative Prices, Journal of Business, 36, 394-419.

Mandelbrot, B. (1971): When Can Price be Arbitraged Efficiently? A Limit to the Validity of the Random Walk and Martingale Model, Review of Economics and Statistics, 53, 225-236.

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

Mantegna, R. and Stanley, E. (2000): An Introduction to Econphysics, Cambridge University Press: Cambridge.

Meese, R. and Rogoff, K. (1983): Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample? Journal Of International Economics, 14, 3-24.

Menkhoff, L. and Schlumberger, M. (1995): Persistent Profitability of Technical Analysis on Foreign Exchange Markets? Banca Nazionale del Lavoro Quarterly Review, 48, 189-216.

Menkhoff, L. (1997): Examining the Use of Technical Currency Analysis, International Journal of Finance and Economics, 2, 307-318.

Menkhoff, L. (1998): The Noise Trading Approach - Questionnaire Evidence from Foreign Exchange, Journal of International Money and Finance, 17, 547-564.

Moosa, I. (2000): Exchange Rate Forecasting: Techniques and Applications, Macmillan: Basingstoke.

Murphy, J. (1999): Technical Analysis of Financial Markets, New York 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.

Neely, C., Weller, P. and Dittmar, R. (1997): Is Technical Analysis in the Foreign Exchange Market Profitable? A Genetic Programming Approach, Journal of Financial and Quantitative Analysis, 32, 405-426.

Neely, C. (1998): Technical Analysis and the Profitability of U.S. Foreign Exchange Intervention, Federal Reserve Bank of St. Louis Review, 84, 3-17.

Obstfeld, M., Taylor, A. (1997): Nonlinear Aspects of Goods-Market Arbitrage and Adjustment: Heckscher's Commodity Points Revisted, Journal of the Japanese and International Economies, 11, 441-479.

Osler C. (1998): Identifying Noise Traders: The Head-And-Shoulders Pattern in US Equities, Federal Reserve Bank of New York Staff Reports, Number 42.

Osler C. (2000): Support for Resistance: Technical Analysis and Intraday Exchange Rates, Federal Reserve Bank of New York Economic Policy Review, 6, 53-68.

Osler C. (2001): Currency Orders and Exchange Rate Dynamics: Explaining the Success of Technical Analysis, Federal Reserve Bank of New York Staff Reports, Number 125.

Ott, E., Sauer, T. and Yorke, J. (1994): Coping with Chaos, Wiley: New York.

Pagan, A. (1996): The Econometrics of Financial Markets, Journal of Empirical Finance, 3, 15-102.

Palmer, R., Arthur, B., Holland, J., LeBaron, B. and Tayler, P. (1994): Artificial Economic Life: A Simple Model of a Stock Market, Physica D, 75, 264-274.

Piron, R. (1991): Keynes as Noise Trader, Journal of Economic Perspectives, 5, 215-217.

Pring, M. (1991): Technical Analysis Explained, McGraw-Hill: New York.

Reiss, R. and Thomas, M. (2000): Statistical Analysis of Extreme Values, Second Edition, Birkhäuser: Basel.

Saacke, P. (1999): Technical Analysis, Central Bank Interventions and Exchange Rate Regimes, PhD Dissertation, University of Hamburg.

Scheinkman, J. and LeBaron, B. (1989): Nonlinear Dynamics and Stock Returns, Journal of Business, 62, 311-337.

Schulmeister, S. (1988): Currency Speculation and Dollar Fluctuations, Banca Nazionale del Lavoro Quarterly Review, 167, 343-365.

Shiller, R. (1990): Speculative Prices and Popular Models, Journal of Economic Perspectives, 4, 55-65.

Shiller, R. (1995): Conversation, Information, and Herd Behavior, American Economic Review, 85, 181-185.

Shleifer, A. and Summers, L. (1990): The Noise Trader Approach to Finance, Journal of Economic Perspectives, 4, 19-33.

Shleifer, A. and Vishny, R. (1997): The Limits of Arbitrage, Journal of Finance, 52, 35-55.

Simon, H. (1955): A Behavioral Model of Rational Choice, Quarterly Journal of Economics, 9, 99-118.

Skouras, S. (2001): Financial Returns and Efficiency as Seen by an Artificial Technical Analyst, Journal of Economic Dynamics and Control, 25, 213-244.

Slovic, P., Lichtenstein, S. and Fischhoff, B. (1988): Decision Making, in: Atkinson, R., Herrnstein, R. Lindzey, G. and Luce, D. (Eds.): Steven's Handbook of Experimental Psychology, Wiley: New York, 673-738.

Sonnemans, J., Hommes, C., Tuinstra, J. and Velden, H. van de (1999): The Instability of a Heterogeneous Cobweb Economy: a Strategy Experiment on Expectation Formation, University of Amsterdam, CeNDEF Working Paper 99-06.

Sprott, J. and Rowlands, G. (1995a): Chaos Data Analyzer: The Professional Version 2.1, American Institute of Physics: New York.

Sprott, J. and Rowlands, G. (1995b): PC User's Manual of Chaos Data Analyzer: The Professional Version 2.1, American Institute of Physics: New York.

Sweeney, R. (1986): Beating the Foreign Exchange Market, Journal of Finance, 41, 163- 182.

Takagi, S. (1991): Exchange Rate Expectations: A Survey of Survey Studies, IMF Staff Papers, 38, 156-183.

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.

Taylor, M. (1995): The Economics of Exchange Rates, Journal of Economic Literature, 33, 13-47.

Tesfatsion, L. (2001): Agent-Based Computational Economics: A Brief Guide to the Literature, in: Michie, J. (Ed.): Reader's Guide to the Social Sciences, Volume 1, Fitzroy-Dearborn: London.

Tobin, J. (1978): A Proposal for Monetary Reform, Eastern Economic Journal, 4, 153-159.

Tversky, A. and Kahneman, D. (1974): Judgment Under Uncertainty: Heuristics and Biases, Science, 185, 1124-1131.

Vries, C. de (1994): Stylized Facts of Nominal Exchange Rate Returns, in: van der Ploeg, F. (Ed.): The Handbook of International Macroeconomics, Blackwell: Oxford, 348-389.

Westerhoff, F. (1999): Japan in der Krise, University of Osnabrück, Japan Research Center, Discussion Paper No. 9, a shorter version is reprinted in: Wacker, G. (Ed.): Asienkrise, Demokratie, Nationalismus. Neue Wechselwirkungen zwischen Politik und Ökonomie in Ostasien, Bundesinstitut für ostwissenschaftliche und internationale Studien: Köln, 23-28.

Westerhoff, F. (2001): Speculative Behavior, Exchange Rate Volatility, and Central Bank Intervention, Central European Journal of Operations Research, in press.

Wolf, A., Swift, J. Swinney, H. and Vastano, J. (1985): Determining Lyapunov Exponents From a Time Series, Physica D, 16, 285-317.

Youssefmir, M. and Huberman, B, (1997): Clustered Volatility in Multiagent Dynamics, Journal of Economic Behavior and Organization, 32, 101-118.

Zeeman, E. (1974): On the Unstable Behavior of Stock Exchanges, Journal of Mathematical Economics, 1, 39-49.