Available online at www.sciencedirect.com
SCIENCE DIRECT®
ELSEVIER Journal of Economic Dynamics & Control 29 (2005) 691–719
JOURNAL OF Economic Dynamics & Control
www.elsevier.com/locate/econbase

Heterogeneity of agents, transactions costs and the exchange rate

Paul De Grauwe[a,*], Marianna Grimaldi[b]

[a]*Department of Economics, University of Leuven, Naamsestraat 69, Leuven 3000, Belgium*
[b]*Sveriges Riksbank, SE 103 37 Stockholm, Sweden*

Received 31 October 2002; accepted 5 January 2004
Available online 7 July 2004

Abstract

We develop a model of the exchange rate that has two features. First, there are non-linearities that arise from the existence of transaction costs in goods markets. Second, the model assumes heterogeneous agents who use simple forecasting rules, the 'fitness' of which is then controlled ex post by checking their profitability, and by switching to the more profitable rules. This model is capable of reproducing the empirical puzzles observed in exchange markets (disconnect puzzle, excess volatility, fat tails, volatility clustering). We analyse some policy implications of this type of modelling of the exchange rate.

JEL classification: F31; F41

Keywords: Exchange rate; Heterogeneous agents; Transaction costs; Chaos; Sensitivity to initial conditions

1. Introduction

Traditional exchange rate modelling has been based on the efficient market rational expectations paradigm. It is increasingly evident, however, that this model is rejected by the data. There is a whole list of empirical puzzles that the traditional

*Corresponding author. Tel.: +32-16 326 794; fax: +32-16 326 796.
E-mail address: paul.degrauwe@econ.kuleuven.ac.be (P. De Grauwe).

0165-1889/$ - see front matter
doi:10.1016/j.jedc.2004.01.004

model fails to explain. The first and foremost empirical puzzle has been called the 'disconnect' puzzle, i.e. the exchange rate appears to be disconnected from its underlying fundamentals most of the time. Goodhart (1989), Goodhart and Figlioli (1991) and more recently Faust et al. (2003) found that most of the changes in the exchange rates occur when there is no observable news in the fundamental economic variables. This finding contradicts the efficient market rational expectations models, which imply that the exchange rate can only move when there is news in the fundamentals.

The exchange rate disconnect puzzle was also implicit in the celebrated Meese and Rogoff studies of the early 1980s (Meese and Rogoff, 1983) documenting that there is no stable relationship between exchange rate movements and the *news* in the fundamental variables.

Other empirical anomalies have been uncovered over the years. One anomaly relates to the existence of excess volatility (Baxter and Stockman, 1989; Flood and Rose, 1995). Other puzzles are that the distribution of the exchange rate returns exhibits fat tails and volatility clustering (see de Vries, 2001; Lux, 1998; Lux and Marchesi, 2000). These empirical anomalies have also been observed in other financial markets (see e.g. Hommes, 2001). This evidence is difficult to rationalise in existing exchange rate models, since there is little evidence of fat tails and volatility clustering in the fundamental variables that drive the exchange rate in these models.

There is a need for other modelling approaches of the exchange rate. Our modelling approach combines two insights. The first one focuses on the presence of non-linearities that arise from the existence of transaction costs in goods markets. Recent research has stressed the importance of transaction costs in the goods market for our understanding of the dynamics of exchange rate adjustments (Obstfeld and Rogoff, 2000; Engel, 2000; Michael et al., 1997; Kilian and Taylor, 2001; Sarno and Taylor, 2002).

The second insight highlights the role of the heterogeneity of agents, who use incomplete information and who have different beliefs about the future exchange rate.[1] Recently heterogeneity of agents was also introduced in rational expectations models (see e.g. Bacchetta and van Wincoop, 2003). The implication of rational expectations in models with heterogeneous agents is that it creates 'infinite regress', i.e. the exchange rate depends on the expectations of other agents' expectations, which depends on the expectations of the expectations of other agents' expectations, and so on, ad infinitum. This leads to intractable mathematical problems except under very restrictive simplifying assumptions. Although this approach is intellectually satisfying, it is unclear that it is a good representation of what agents do in the exchange market. It requires these agents to solve a mathematical problem to which mathematicians have as yet been unable to give a general solution. This seems to us as imposing too large an informational burden on individual agents.

[1]It should be noted that the heterogeneity of agents' expectations has been recognised as being important to explain the dynamics of asset prices, including the exchange rate (see De Long et al., 1990; Frankel and Froot, 1986; Brock and Hommes, 1998; Lux and Marchesi, 2000; Hommes, 2001).

Our approach contrasts with this rational expectations approach in that agents use simple rules, the 'fitness' of which is then controlled ex post by checking their profitability,[2] and by switching to the more profitable rules.

The paper is organised as follows. In Section 2 and 3 we present the theoretical model. In Sections 4–7 we analyse its features, while in Section 8 we show the empirical relevance of the model. We conclude with some general implications for the exchange rates of the major currencies.

2. A simple non-linear exchange rate model

In this section we develop a simple non-linear exchange rate model. We start by defining the fundamental exchange rate. This is the exchange rate that is consistent with equilibrium in the real part of the economy. In a very simple model this could be the PPP-value of the exchange rate. In more elaborate models (e.g. the monetary model or the Obstfeld–Rogoff new open economy macro model (Obstfeld and Rogoff, 1996)) this fundamental exchange rate could be determined by the interaction of more variables than the price levels. We leave the modelling of the fundamental exchange rate outside the scope of this paper, and we will assume that the fundamental exchange rate behaves like a random walk without drift.[3] This implies

$$s_t^* = s_{t-1}^* + \varepsilon_t. \tag{1}$$

We now introduce the assumption that the agents have heterogeneous beliefs. We assume two types of agents, which we will call *fundamentalists* and *chartists*.[4]

The *fundamentalists* compare the current market exchange rates with the fundamental rate and they forecast the future market rate to move towards the fundamental rate. In this sense they follow a *negative feedback* rule.[5] We will make the additional assumption that they expect the speed with which the market rate returns to the fundamental rate to be determined by the speed of adjustment in the goods market, θ which we assume to be constant. This leads us to specify the following rule for the fundamentalists:

$$E_{f,t}(\Delta s_{t+1}) = -\theta(s_t - s_t^*), \tag{2}$$

where $E_{f,t}$ is the forecast made in period t by the fundamentalists using information up to time t, s_t is the exchange rate in period t, and $\theta > 0$.

[2]See Brock and Hommes (1998) for an application to the stock market. By stressing the use of simple rules, this approach comes close to the one of behavioural finance (Schleiffer, 2000).

[3]Introducing a drift does not change the nature of the model, nor its results. We also experimented with an AR(1) process for the fundamental rate. This did not affect our results.

[4]This way of modelling the foreign exchange market was first proposed by Frankel and Froot (1986). It was further extended by De Long et al. (1990) and De Grauwe et al. (1993) and more recently Kilian and Taylor (2001). For evidence about the use of chartism see Taylor and Allen (1992).

[5]Note that this is also the approach taken in the Dornbusch model.

The *chartists* are assumed to follow a *positive feedback* rule, i.e. they extrapolate past movements of the exchange rate into the future. Their forecast is written as

$$E_{c,t}(\Delta s_{t+1}) = \beta \sum_{i=0}^{T} \alpha_i \Delta s_{t-i}, \tag{3}$$

where $E_{c,\,t}$ is the forecast made by the chartists using information up to time t, Δs_t is the change in the exchange rate, $\sum_{i=0}^{T} \alpha_i = 1$, and $0 < \beta < 1$ to ensure dynamic stability.

As can be seen, the chartists compute a moving average of the past exchange rate changes and they extrapolate this into the future exchange rate change. The degree of extrapolation is given by the parameter β. Note that the chartists do not take into account information concerning the fundamental exchange rate. In this sense they can be considered to be pure *noise traders* (see De Long et al., 1990).

Our choice to introduce chartists' rules of forecasting is based on empirical evidence. The evidence that chartism is used widely to make forecasts is overwhelming (see Cheung et al., 1999; Taylor and Allen, 1992). Therefore, we give a prominent role to chartists in our model. It remains important, however, to check if the model is internally consistent. In particular, the chartists' forecasting rule must be shown to be profitable within the confines of the model. If these rules turn out to be unprofitable, they will not continue to be used.

The next step in our analysis, therefore, is to specify how agents evaluate the usefulness of these two forecasting rules. The general idea that we will follow is that agents use one of the two rules, compare their profitability ex post and then decide whether to keep the rule or switch to the other one. Thus, our model is in the logic of evolutionary dynamics, in which simple decision rules are followed. These rules continue to be followed if they pass some 'fitness' test (profitability test).

In order to implement this idea we follow the procedure proposed by Brock and Hommes (1997, 1998) which is based on discrete choice theory. Thus, the fractions of agents using chartist and fundamentalist rules are a function of the relative (risk adjusted) profitability of these rules, i.e.

$$n_{ct} = \frac{\exp\, \gamma(\pi_{c,t-1} - \sigma_{c,t-1})}{\exp\, \gamma(\pi_{c,t-1} - \sigma_{c,t-1}) + \exp\, \gamma(\pi_{f,t-1} - \sigma_{f,t-1})}, \tag{4}$$

$$n_{ft} = \frac{\exp\, \gamma(\pi_{f,t-1} - \sigma_{f,t-1})}{\exp\, \gamma(\pi_{c,t-1} - \sigma_{c,t-1}) + \exp\, \gamma(\pi_{f,t-1} - \sigma_{f,t-1})}, \tag{5}$$

where $\pi_{c,t-1}$ and $\pi_{f,t-1}$ are the net profits of the chartists' and fundamentalists' forecasting rule in period $t-1$. Chartists and fundamentalists make a profit (loss) when they correctly (wrongly) forecast the direction of the exchange rate movements. The profit (loss) they make equals the one-period return of the exchange rate. We assume that fundamentalists bear a fixed cost for collecting information on the fundamental variable while chartists' information is costless. Thus

$$\pi_{f,t-1} = \pi'_{f,t-1} - C_f,$$

where $\pi\prime_{f,t-1}$ *is the* gross profit of fundamentalists and C_f is the fixed cost of collecting information on the fundamental.

The variables $\sigma_{c,t-1}$ and $\sigma_{f,t-1}$ are measures of the risk associated with the forecasting rule of the chartists and fundamentalists, respectively. The most obvious definition of these measures of risk is the weighted average of the squared (one period ahead) forecast errors made by chartists and fundamentalists, respectively:

$$\sigma_{i,t} = \sum_{k=1}^{\infty} \lambda_k [E_{i,t-k}(s_{t-k+1}) - s_{t-k+1}]^2, \tag{6}$$

where λ_k are geometrically declining weights.

This definition of the risk leads to a problem for the fundamentalists, however. The latter make a forecast based on the distance between the market exchange rate and the fundamental rate (the misalignment). As a result, when the exchange rate departs from its fundamental, the squared forecast error of using a fundamentalist rule increases. Thus, if we use (6) as a measure of risk for the fundamentalists, it implies that the stronger the misalignment, the riskier the use of a fundamentalist forecast will be perceived to be. This is quite implausible. One would expect that as the degree of misalignment increases, the confidence in making forecasts based on a fundamentalist rule also increases. In order to take this feature of fundamentalists forecasts into account, we amend Eq. (6) for the fundamentalists as follows:

$$\sigma_{f,t} = \frac{\sum_{k=1}^{\infty} \lambda_k [E_{i,t-k}(s_{t-k+1}) - s_{t-k+1}]^2}{1 + (s_{t-1} - s_{t-1}^*)^2}, \tag{7}$$

where $(s_{t-1} - s_{t-1}^*)$ is the misalignment. The logic behind this specification is that with increasing misalignment the fundamentalists attach less importance to the short-term volatility as measured by the one period ahead forecast error, and they become increasingly confident that the exchange rate will revert to its fundamental value.[6] Note that in the neighbourhood of the fundamental exchange rate the risk variable converges to the same value as in (6).

Eqs. (4) and (5) can now be interpreted as follows. When the risk adjusted profits[7] of the chartists rule increase relative to the risk adjusted profits of the fundamentalists rule, then the fraction of the chartists in the market increases, and vice versa. The sensitivity with which the chartists' and fundamentalists' fractions adjust to the relative profitability of the forecasting rules depends on the parameter γ. With an increasing γ the fraction of chartists (fundamentalists) who switch to the more profitable forecasting rule increases. In the limit when γ goes to infinity agents will select the most profitable rule instantaneously. When γ is equal to zero the fraction of chartists and fundamentalists is constant and equal to 0.5. Thus γ is a measure of inertia in the decision to switch to the more profitable rule.[8]

[6] For a similar approach see Chiarella et al. (2002).

[7] Note that the risk adjusted profits can be interpreted as Sharpe ratios.

[8] This specification of the decision rule is often used in discrete choice models. See for example Brock and Hommes (1997) and Lux (1998).

The market expectation of the exchange rate change can be written as a weighted average of the expectations of chartists and fundamentalists, i.e.

$$E_t\Delta s_{t+1} = -n_{ft}\theta(s_t - s_t^*) + n_{ct}\beta\sum_{i=0}^{T} a_i\Delta s_{t-i}, \tag{8}$$

where n_{ft} and n_{ct} are the weights of fundamentalists and chartists respectively.

The realised change in the market exchange rate in period $t+1$ equals the market forecast made at time t plus some white noise errors (i.e. the *news* that could not be predicted at time t):

$$\Delta s_{t+1} = -n_{ft}\theta(s_t - s_t^*) + n_{ct}\beta\sum_{i=0}^{T} a_i\Delta s_{t-i} + \varepsilon_{t+1}. \tag{9}$$

3. The model with transaction costs

There is an increasing body of theoretical literature stressing the importance of transaction costs in the goods market as a source of non-linearity in the determination of the exchange rate (Dumas, 1992; Obstfeld and Rogoff, 2000). The importance of transaction costs in the goods market has also been confirmed empirically (Taylor et al., 2001; Kilian and Taylor, 2001). It should be noted that transaction costs in the goods market remain sizeable because a large component of most tradable goods is non-tradable (see Obstfeld and Rogoff, 2000).

We therefore introduce transaction costs into the model and we assume that the fundamentalists take the existence of transaction costs in the goods market into account, i.e. they behave differently depending on whether the exchange rate is within or outside the transaction costs band. When the exchange rate deviations from the fundamental value are smaller than the transaction costs in the goods markets, there is no mechanism that drives the exchange rate towards its equilibrium value. As a result, fundamentalists expect the changes in the exchange rate to follow a white noise process ε_t. The best they can do is to forecast no change. More formally,

$$\text{when } |s_t - s_t^*| < C, \quad \text{then } E_{f,t}(\Delta s_{t+1}) = 0.$$

In the second case, when the exchange rate deviation from its fundamental value is larger than the transaction costs C (assumed to be of the 'iceberg' type), then the fundamentalists follow the same forecasting rule as in Eq. (2). More formally,

$$\text{when } |s_t - s_t^*| > C \text{ holds, then Eq. (2) applies.}$$

This formulation implies that when the exchange rate moves outside the transaction costs band, market inefficiencies other than transaction costs continue to play a role. As a result, these inefficiencies prevent the exchange rate from adjusting

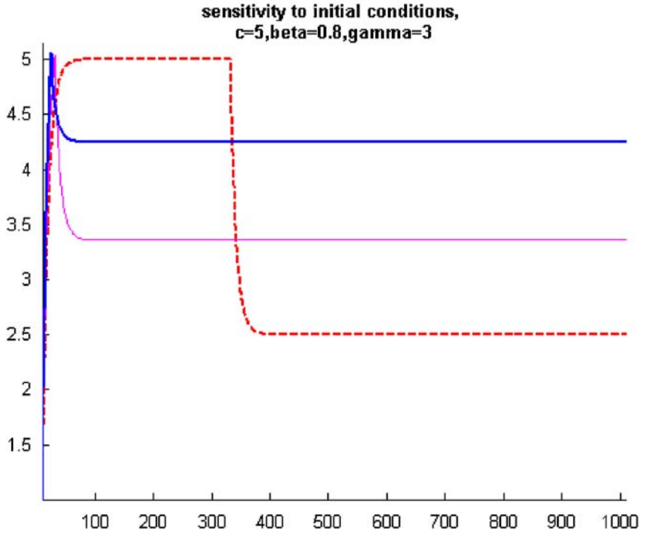

Fig. 1. *Note*: The parameters α_i and λ_k were defined as geometrically declining weights with $\rho = 0.6$; $C_f = 0.05$. These parameter values were kept unchanged in all the simulations reported here. A sensitivity analysis revealed that the results are not very sensitive to these parameter values.

instantaneously. In our model these inefficiencies are captured by the fact that the speed of adjustment in the goods market is not infinite (Eq. (2)).

4. Solution of the model

In this section we investigate the properties of the solution of the model. We first analyse the deterministic part of the model so as to obtain a better insight into the characteristics of the solution that is not clouded by exogenous noise. We use simulation techniques since the non-linearities do not allow for a simple analytical solution. We select 'reasonable' values of the parameters, i.e. those that come close to empirically observed values. As we will show later these are also parameter values for which the model replicates the observed statistical properties of exchange rate movements. We will however analyse how sensitive the solution is to different sets of parameter values.

We first concentrate on the fixed-point solutions of the model. We find that for a relatively wide range of parameters the solution converges to a fixed point (a *fixed-point attractor*). However, there are many such fixed points (attractors) to which the solution converges depending on the initial conditions.[9] We illustrate this feature in Fig. 1, where we show the exchange rate in the time domain for a particular set of parameters and different initial conditions. We find that the exchange rate converges to a different fixed point depending on the initial conditions. (In the next section we

[9]Note that the initial condition that is changed refers to the one period lagged exchange rate.

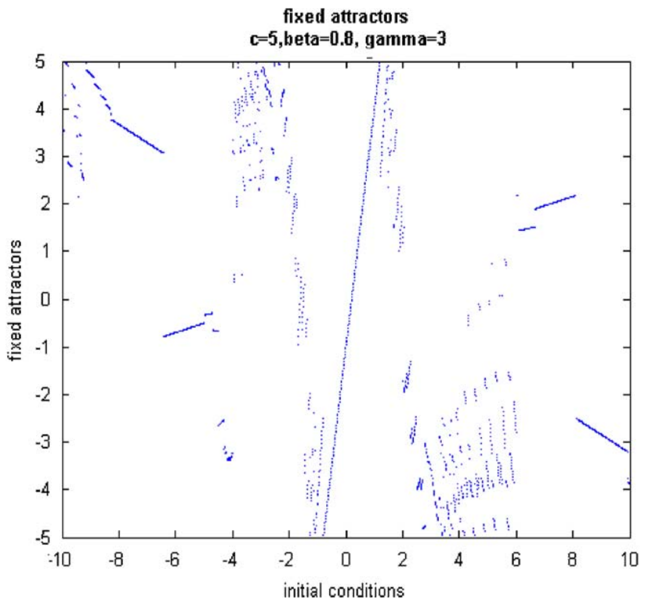

Fig. 2. *Note*: The parameters α_i and λ_k were defined as geometrically declining weights with $\rho = 0.6$; $C_f = 0.05$. These parameter values were kept unchanged in all the simulations reported here. A sensitivity analysis revealed that the results are not very sensitive to these parameter values.

perform a sensitivity analysis to check the general nature of this result.) We show this feature in Fig. 2 in a different way by plotting the fixed-point solutions (attractors) as a function of the different initial conditions. On the horizontal axis we set out the different initial conditions. These are initial shocks to the deterministic system. The vertical axis shows the solutions corresponding to these different initial conditions. Note that small changes in the initial conditions lead to large and discontinuous displacements of the attractors. This characteristic is a natural result of the non-linear nature of our model. We return to this to give an interpretation of this phenomenon.

5. Sensitivity analysis

We obtain a multiplicity of fixed-point solutions for a relatively broad range of parameters. We find that the extrapolation parameter of the chartists, β and the intensity of choice parameter γ are of crucial importance. In Fig. 3 we show the fixed-point attractors for different combinations of parameter values of β and γ.

It can be seen that we obtain a multiplicity of fixed-point attractors, each one depending on the initial shock. It should also be noted that the fixed-point attractors lie within the transaction costs band. The intuition is that any fixed-point solution outside the transaction costs band would create an inconsistency, which can be described as follows. Outside the transaction costs band the fundamentalists'

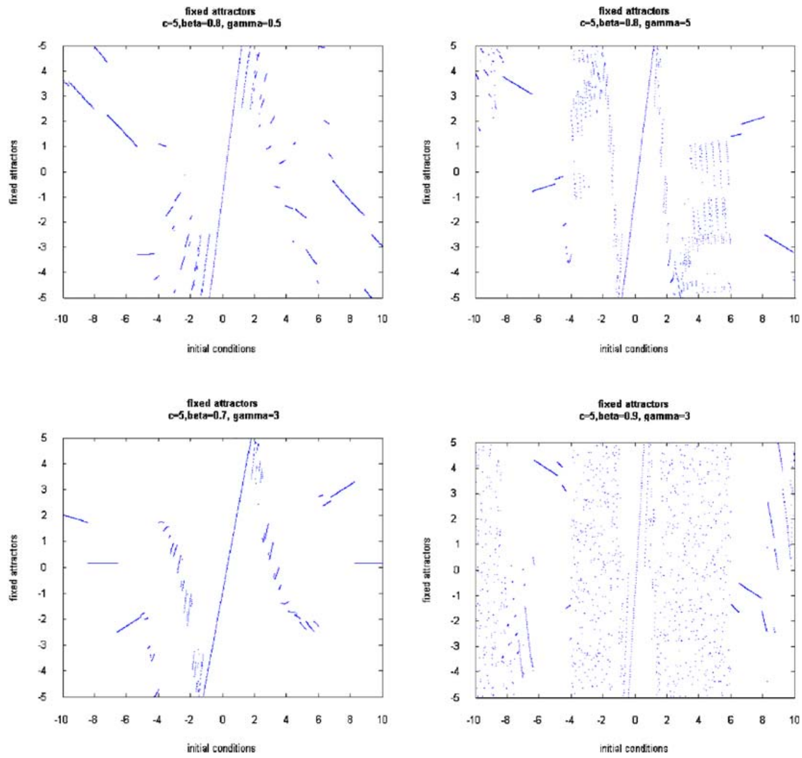

Fig. 3.

behaviour leads to a mean reverting process of the exchange rate, moving the latter towards the transaction costs band. Thus, if a fixed-point solution were observed outside the transactions cost band, this would mean that the fundamentalists would fail to move the exchange rate towards the band. Once inside the band, the fundamentalists' dynamics disappears. The only dynamics then comes from the chartists who drive the exchange rate to some attractor within the band. The exact position of this attractor depends on the entry point of the exchange rate in the transactions cost band, and this depends on the initial shock.

We next perform a more extensive sensitivity analysis. We analyse how changes in some important parameters of the model affect the nature of the solution. We first do a sensitivity analysis with respect to the transaction cost parameter C. In Fig. 4 we show the attractors as a function of transaction costs. We observe that as transaction costs increase the band in which the fixed-point attractors are located increases correspondingly. The important aspect of Fig. 4 is that small changes in the transaction costs lead to a large and discontinuous displacement of the fixed attractor. This feature was also found when we plotted the fixed attractors for different initial conditions (see Figs. 2 and 3).

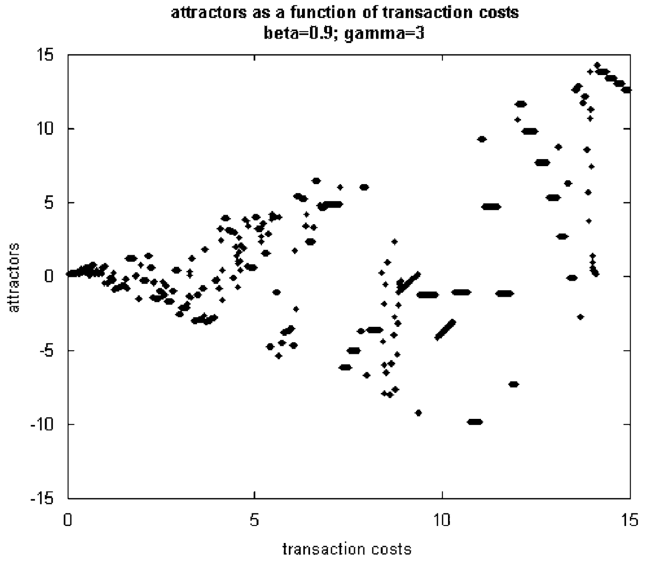

Fig. 4.

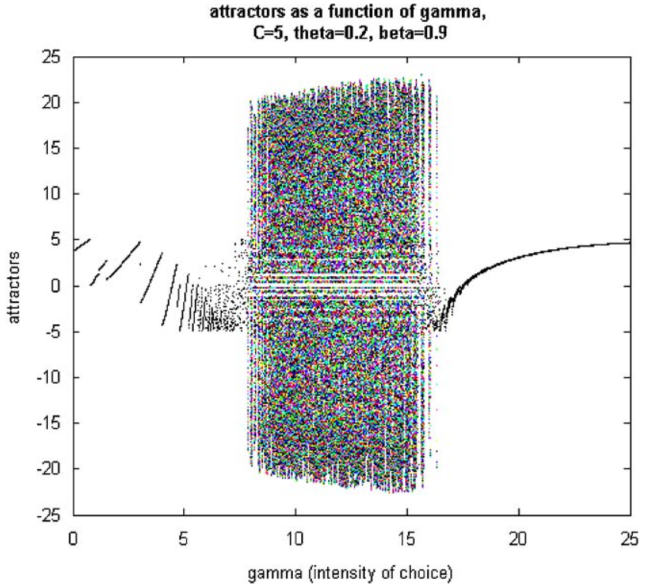

Fig. 5.

Next we perform a similar sensitivity analysis by allowing changes in γ, the intensity of choice parameter. In Fig. 5 we show the equilibrium exchange rate (attractor) as a function of γ. We observe that for relatively low values of γ we obtain fixed-point solutions. For intermediate values of γ we obtain a chaotic region, i.e. the

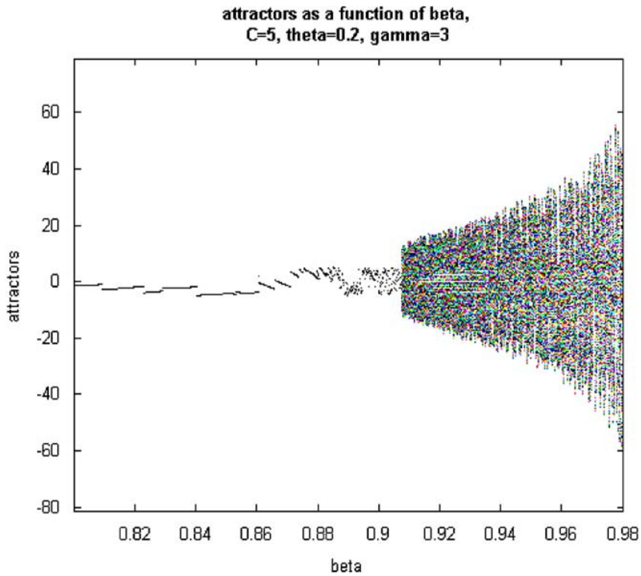

Fig. 6.

exchange rate moves within a strange attractor. For high values of γ we return to fixed-point solutions.[10] These, however, are 'well-behaved', i.e. they lie on a continuous line. This means that small changes in the parameter γ stop having discontinuous effects on the equilibrium exchange rate. Thus, when agents react forcefully to the relative profitability of the forecasting rules the system converges to a unique fixed-point solution.

Finally in Fig. 6 we show the sensitivity of the equilibrium exchange rate with respect to the extrapolation parameter of chartists, β. For values of $\beta < 0.95$ we obtain fixed-point attractors. In this region we obtain the same characteristic that we observed in the sensitivity analysis with respect to the transaction cost, i.e. a small change in the parameter leads to discrete jumps in the equilibrium exchange rate. When β exceeds 0.91 we enter a chaotic region.

The empirical evidence about the existence of deterministic chaos in the foreign exchange rate market is weak (see Guillaume, 1996; Schittenkopf et al., 2001). Therefore, we will focus the analysis of the model on parameter values that do not lead to deterministic chaos. We will show that in combination with stochastic shocks this model is capable of producing a dynamics that exhibits many of the features of chaotic dynamics despite the fact that the deterministic solutions of the model are fixed points.

[10]It should be pointed out that the numerical values of γ that move us in and out of chaos depend on the initial conditions of the simulation runs, and on the other parameters of the model. In all simulations, however, we obtain three regions: a multitude of fixed points for low values of γ, chaotic attractors for intermediate values of γ, a unique fixed point for high values of γ.

6. The stochastic version of the model

We now introduce stochastic disturbances to the model. These disturbances affect the fundamental, which is assumed to be a random walk. In addition, as can be seen from Eq. (9), there is exogenous noise leading to forecast errors of chartists and fundamentalists. We simulate the model with a certain combination of parameter values that we refer to as the 'standard case'. This includes setting $c = 5$, $\beta = 0.9$, $\theta = 0.2$ and $\gamma = 3$. Our results hold for a wide range of parameter values. As mentioned earlier these are parameter values that do not produce deterministic chaos.

A first feature of the solution of the stochastic version of the model is the sensitivity to initial conditions. In order to show this, we first simulated the model with the 'standard' parameter values and then with the same parameters setting but with a slightly different initial condition. In both cases we used identical stochastic disturbances. We show the time paths of the (market) exchange rate in Fig. 7. We observe that after a certain number of periods the two exchange rates start following a different path. This result is related to the presence of many fixed-point attractors in the deterministic part of the model, which are themselves dependent on the initial conditions (see Fig. 3, which shows how slight differences in initial conditions can lead to fixed-point attractors that are very far apart). As a result, the two exchange rates can substantially diverge because attracted by fixed-points that are located in different basins of attraction. The interesting feature of this result is that the combination of exogenous noise and a multiplicity of fixed-point attractors creates chaos-like dynamics without chaos being present in the deterministic part of the model.

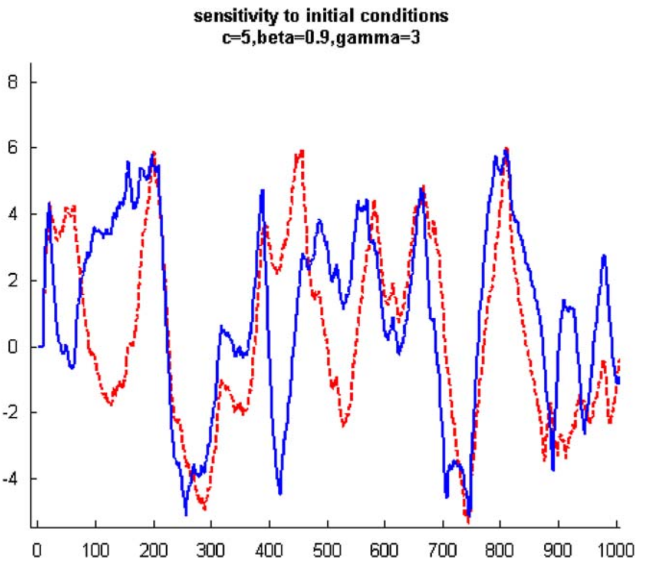

Fig. 7.

7. The effect of permanent shocks

In this section we analyse how a permanent shock in the fundamental exchange rate affects the market exchange rate. In linear models a permanent shock in the fundamental has a predictable effect on the exchange rate, i.e. the coefficient that measures the effect of the shock in the fundamental on the exchange rate converges after some time to a fixed number. Things are very different in our non-linear model. We illustrate this by showing how a permanent increase in the fundamental is transmitted to the exchange rate. We assumed that the fundamental rate increases by 10, and we computed the effect on the exchange rate by taking the difference between the exchange rate with the shock and the exchange rate without the shock. The simulations of these two exchange rates are done using the same exogenous noise. In a linear model we would find that in the long run the exchange rate increases by exactly 10. This is not the case in our model. We present the evidence in Fig. 8 where we show the effect of a permanent shock of 10 in the fundamental rate on the exchange rate for our standard set of parameter values.

The most striking feature of these results is that the effect of the permanent shock does not converge to a fixed number. In fact, it follows a complex pattern. Thus, in a non-linear world it is very difficult to predict what the effect will be of a given shock in the fundamental, even in the long run. Such predictions can only be made in a statistical sense, i.e. our model tells us that *on average* the effect of a shock of 10 in

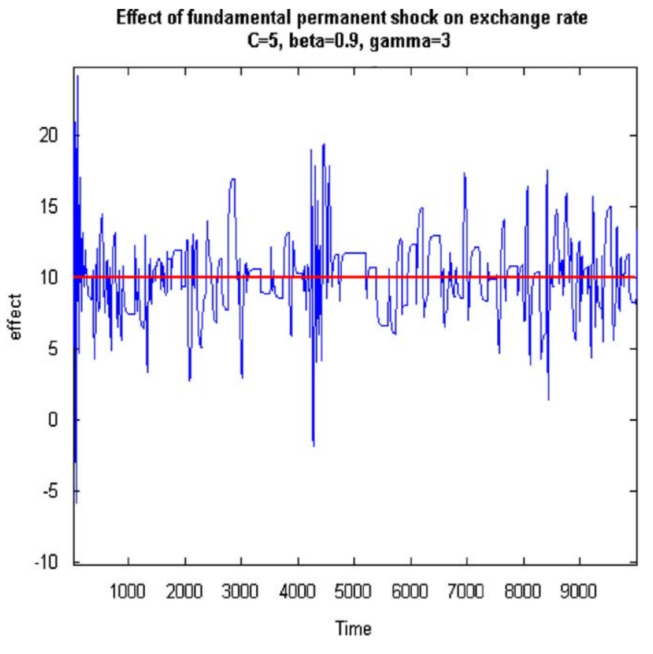

Fig. 8.

the fundamental will be to increase the exchange rate by 10. In any given period, however, the effect could deviate substantially from this average prediction.

The importance of the initial conditions for the effect of a permanent shock in the fundamental can also be seen by the following experiment. We simulated the same permanent shock in the fundamental but applied it in two different time periods. In the first simulation we applied the shock in the first period; in the second simulation we applied it in the next period. The exogenous noise was identical in both simulations. Thus the only difference is in the timing of the shock. We show the results in Fig. 9.

We observe that the small difference in timing changes the future history of the exchange rate. As a result, the effect of the shock measured at a particular point in time can be very different in both simulations. Thus history matters. The time at which the permanent shock occurs influences the effects of the shock.

Note however that in a statistical sense, timing does not matter. When we compute the *average* effect of the same shock in the two simulations over a sufficiently long period of time we obtain the same result, i.e. the exchange rate increases by 10 on average. The time period needed to make valid statistical inferences, however, is large. We illustrate this in Fig. 10 by the frequency distributions of the effects of the same shock in the two simulations obtained over two different simulation runs, the first one containing 10 000 periods, the second one 1000 periods.

An important aspect of Fig. 10 is that when computed over a sample of 1000 periods, the distribution of the effects is irregular and quite different for the two simulations. Only when the sample becomes very large (10 000) do we obtain

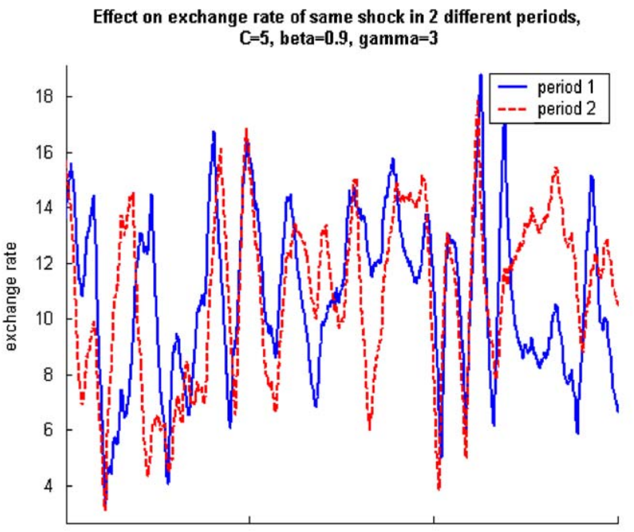

Fig. 9.

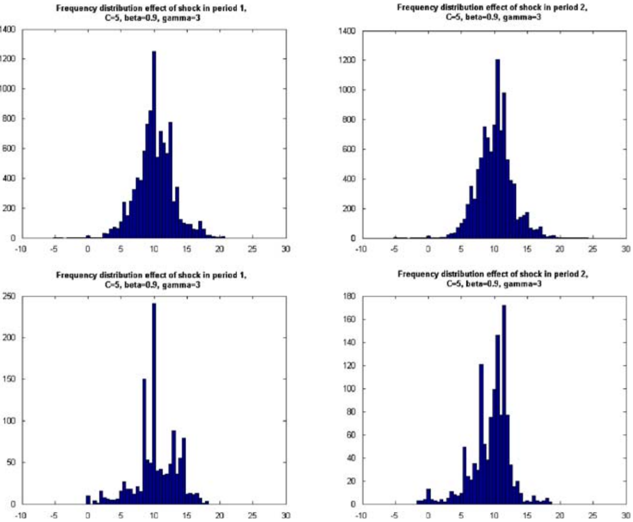

Fig. 10.

'well-behaved' distributions permitting statistical inferences about the effect of the same shock.[11]

Our results help to explain why in the real world it appears difficult to predict the effects of changes in the fundamental exchange rate on the market rate, and why these effects seem to be very different when applied in different periods. In fact this is probably one of the most intriguing empirical problems. Economists usually explain the difficulty of forecasting the effects of a particular change in one exogenous variable (e.g. an expansion of the money stock) by invoking the ceteris paribus hypothesis, i.e. there are usually other exogenous variables changing unexpectedly, preventing us to isolate the effect of the first exogenous variable. In our model the uncertainty surrounding the effect of a disturbance in an exogenous variable is not due to the failure of the ceteris paribus hypothesis. No other exogenous variable is allowed to change. The fact is that the change in the exogenous variable occurs at a particular time, which is different from all other times. Initial conditions (history) matter to forecast the effect of shocks. Since each initial condition is unique, it becomes impossible to forecast the effect of a shock at any given point in time with any precision.

[11] We computed tests of normality of the distribution (Kolmogorov–Smirnov test, and Lillie test). We rejected normality in all cases. Thus, the usual significance tests that assume normality of the distribution (t-test and F-test) are not appropriate here.

8. Empirical relevance of the model

In this section we analyse how well our model mimics the empirical anomalies and puzzles that have been uncovered by the flourishing empirical literature. We calibrate the model such that it replicates the observed statistical properties of exchange rate movements. The parameters of the model that do this are those that we used in the previous sections. As was noted there, typically these are parameter sets that do not produce deterministic chaos. We start with the 'disconnect puzzle'.

8.1. The disconnect puzzle

The 'disconnect' puzzle (see Obstfeld and Rogoff, 2000) states that the exchange rate is disconnected from its underlying fundamentals most of the time.[12] It was first analysed by John Williamson (1985) who called it the 'misalignment problem'. This puzzle was also implicit in the celebrated Meese and Rogoff studies of the early 1980s documenting that there is no stable relationship between exchange rate movements and the *news* in the fundamental variables.

Our model is capable of mimicking this empirical regularity. In Fig. 11 we show the market exchange rate and the fundamental rate for a combination of parameters that does not produce deterministic chaos. We observe that the market rate can deviate from the fundamental value substantially and in a persistent way. Moreover, it appears that the exchange rate movements are often disconnected from the movements of the underlying fundamental. In fact, they often move in opposite directions.

We show the nature of the disconnect phenomenon in a more precise way by applying a cointegration analysis to the simulated exchange rate and its fundamental using the same parameter values as in Fig. 11 for a sample of 8000 periods. We found that there is a cointegration relationship between the exchange rate and its fundamental.[13] Note that in our setting there is only one fundamental variable. This implies that no bias from omitted variables can occur.

In the next step we specify an EC model in the following way:

$$\Delta s_t = \mu(s_{t-1} - \gamma s^*_{t-1}) + \sum_{i=1}^{n} \lambda_i \Delta s_{t-i} + \sum_{i=1}^{n} \varphi_i \Delta s^*_{t-i}. \tag{10}$$

The first term on the right-hand side is the error correction term. The result of estimating this equation is presented in Table 1 where we have set $n = 4$.[14]

[12]In its original formulation the disconnect puzzle has two dimensions. One says that the exchange rate is disconnected from its fundamental. The second dimension relates to the fact that real variables (for example, the trade account) do not react to the changes in the exchange rate. In this paper we only analyse the first dimension.

[13]We first performed a unit root test on the simulated exchange rate. We could not reject the existence of unit root. Next, we tested for cointegration using the Johansen cointegration procedure (see Johansen (1991)). We assumed that there is no deterministic trend in the data. However we do allow the intercept different from zero.

[14]The number of lags has been chosen according to the information criteria, such that the error term is white noise.

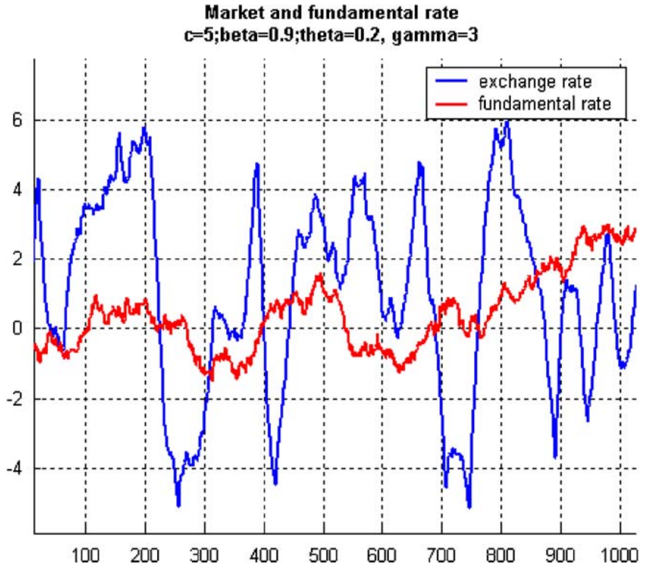

Fig. 11.

Table 1
Parameter estimates of EC model (Eq. (10))

Error correction term		Δs_{t-i}				Δs^*_{t-i}			
μ	γ	λ_1	λ_2	λ_3	λ_4	φ_1	φ_2	φ_3	φ_4
−0.006	1.001	0.53	0.19	0.09	0.06	0.06	0.02	−0.01	−0.01
−15.4	*24.8*	*52.7*	*16.6*	*7.80*	*6.1*	*4.9*	*1.3*	*−1.3*	*−0.4*

Note: The sample consists of 8000 periods. The numbers in italics are t-statistics. $R^2 = 0.63$.

We find that the error correction coefficient (μ) is very low. This suggests that the mean reversion towards the equilibrium exchange rate takes a very long time. In particular, only 0.6% of the adjustment takes place each period. It should be noted that in the simulations we have assumed a speed of adjustment in the goods market equal to 0.2. This implies that each period the adjustment in the goods market is 20%. Thus the nominal exchange rate is considerably slower to adjust towards its equilibrium than what is implied by the speed of adjustment in the goods market. This slow adjustment of the nominal exchange rate is due the chartist extrapolation behaviour.[15] From Table 1, we also note that the changes in fundamentals have a small and insignificant impact on the change in exchange rate. In contrast, the past changes in the exchange rate play a significant role in explaining the change in

[15]Cheung et al. (2001) have provided evidence indicating that the slow speed of adjustment towards PPP is not so much due to the slow speed of adjustment of prices in the goods market but to the slow speed of adjustment of the nominal exchange rate.

exchange rate.[16] These results are consistent with the empirical findings using VAR approach, which suggests that the exchange rate is driven by its own past (see De Boeck, 2000).

We also performed a cointegration analysis for shorter sample periods (1000 periods). We find that in some sample periods the exchange rate and its fundamental are cointegrated, in other sample periods we do not find cointegration. This is in line with the empirical evidence indicating that in some periods the exchange rate seems to be disconnected from its fundamental while in other periods it tightly follows the fundamentals.[17]

The results reported in this section use a particular numerical value for the transactions costs, i.e. $C = 5$. An implication of this choice is that the exchange rate remains within the transaction cost band most of the time. In fact, in the simulation reported in Fig. 11, the exchange rate remains within the transactions cost band 97% of the time. Thus our explanation of the misalignment puzzle relies very much on the existence of a relatively wide band of transactions costs compared to the variability of the underlying fundamentals. Is this a reasonable explanation for the existence of misalignments? Our claim is that it is the right one for the currencies of the industrial countries like the US, Japan and the European countries. There are two pieces of evidence that substantiate this claim.

First, transactions costs in international trade continue to be substantial, as important empirical evidence shows. In particular, several recent empirical studies report the continued existence of large price differentials for the same traded goods across borders (see Haskel and Wolf, 2001; Engel and Rogers, 1995). In Table 2 we provide additional evidence. We show the price dispersion of a sample of identical products in the European Union.[18] We observe that price differentials of up to 40% occur both in the category of foodstuff and of electronic products. This indicates that producers apply ‘pricing to market’. Such pricing strategies, however, can only be applied successfully if transaction costs prevent arbitrage. Thus, the large observed price differentials suggest that transactions costs for traded goods are large and of the order of 20–40%. In addition, for many services, which are non-traded goods, transactions costs are even higher. (See Obstfeld and Rogoff (2000) who argue that transactions costs are key to understanding the major puzzles in international economics.)

Second, the size of the shocks in the fundamentals driving the exchange rates of the major currencies is typically small. These fundamentals include inflation differentials, differentials in interest rates, in growth rates of the money stock, and in growth rates of output. These differentials are typically a few percentage points per year. Thus, one can conclude that the exchange rates of the major currencies

[16]It should be noted that our results are akin to what was found in stock markets, i.e. that in the short-run the exchange rate under-reacts to news, while it overreacts in the long run. See Shleifer (2000).

[17]See Obstfeld and Rogoff (2000). See also De Grauwe and Grimaldi (2001) for a survey of the empirical evidence. In De Grauwe and Vansteenkiste (2001) we present additional empirical evidence.

[18]Similar price differentials exist for other product groups. The large price differentials for cars are notorious. Interestingly, The Economist magazine, which champions the cause of free competition, applies differences in its subscription rates of 20% across the Eurozone (in 2003).

Table 2
Inter-country price dispersion for selected products (excluding VAT), 2000

Supermarket products	
Evian mineral water	43%
Rexona deodorant	21%
Sensodyne toothpaste	21%
Mars bars (single)	21%
Mars bars (multipack)	22%
Coca Cola	21%
Pedigree pal dog food	10%
Plenitude face care	21%
Colgate toothpaste	14%
Bonne Maman marmelade	19%
Electronic products	
Philips audio system	28%
Sony audio system	38%
Canon camcorder	32%
Panasonic portable CD	40%
Philips portable CD	56%
Pioneer CD player	34%
Sony CD player	28%
Philips TV (14 in)	41%
Sony TV (14 in)	33%
Panasonic TV (28 in)	25%
Philips TV (28 in)	61%
JVC VCR	30%
Panasonic VCR	22%
Sony VCR	44%

Source: European Commission, *Price dispersion in the internal market*, and *Price differentials for supermarket goods in the EU*. Both documents can be downloaded from www.europa.eu.int.
Note: Price dispersion is defined as the percentage difference between the most expensive and the cheapest item.

move in an environment in which the shocks in the fundamentals are relatively small compared to the size of transactions costs. As a result, the exchange rates of these currencies move most of the time within a band within which few opportunities exist for goods market arbitrage. This considerably weakens the mean reversion dynamics on which fundamentalism is based.

It is important to analyse the dynamics of the exchange rate under different combinations of transactions costs and size of shocks in fundamentals. After all, there are many countries in the world where the size of the shocks in fundamentals is very large compared to transactions costs (e.g. Latin American countries that have experienced triple digit inflation rates and growth rates of their money stocks). The way we proceed is to simulate the model under different assumptions about the size of transactions costs, while keeping the size of the shocks unchanged. (Note that we could also vary the size of the fundamental shocks while keeping the transactions cost unchanged. This gives qualitatively the same results.)[19] We then apply the same

[19]In Section 9 we show such an exercise.

Table 3
Transaction costs and speed of adjustment

Transaction cost (C)	Error correction coeff. (μ)	Percent of time outside transaction cost band (%)
0.5	−0.062	44
1.0	−0.044	25
2.5	−0.016	7
5.0	−0.006	3

error correction model on the estimated exchange rate as in Eq. (10). We show the coefficients of the error correction term (which is a measure of the speed with which the exchange rate returns to its fundamental value) in Table 3. We observe that there is an inverse relationship between the size of transaction costs and the speed of adjustment. With low transactions costs (relative to the size of fundamental shocks) the speed of adjustment is high; with high transaction costs the speed of adjustment is low. Thus in a world where the transaction costs are small relative to the size of the fundamental shocks misalignments are quickly corrected. In such a world the exchange rate is pushed outside the transactions cost band frequently (see last column of Table 3) so that the mean reverting forces originating from goods market arbitrage are forceful.

Thus, our model generates an empirical regularity (the 'disconnect' puzzle) that has also been observed in reality. We can summarise the features of this puzzle as follows. First, over the very long run the exchange rate and its fundamentals are cointegrated. However, the speed with which the exchange rate reverts to its equilibrium value is very slow. Second, in the short run the exchange rate and its fundamentals are 'disconnected', i.e. they do not appear to be cointegrated. Third, the nature of the disconnect puzzle changes depending on the relative size of transactions costs versus the size in the fundamental shocks. When the size of fundamental shocks is small relative to transactions costs misalignment is relatively long and protracted. This is the case with the currencies of the major industrial countries. When the size of the fundamental shocks is large relative to transactions costs, misalignments although large are quickly corrected. The empirical evidence substantiates these results (see Sarno and Taylor (2002) who show that when the size of the shocks to the PPP-relation is large, the speed of adjustment towards PPP is also high).

8.2. Fat tails and excess kurtosis

It is well known that the exchange rate changes do not follow a normal distribution. Instead it has been observed that the distribution of exchange rate changes has more density around the mean than the normal and exhibits fatter tails than the normal (see de Vries, 2001). This phenomenon was first discovered by Mandelbrot (1963) in commodity markets. Since then, fat tails and excess kurtosis

Table 4
Measure of fat tails: the Hill index

Parameter values	Kurtosis	Median Hill index (5 samples 2000 observations)		
		2.5% tail	5% tail	10% tail
$C = 5$, $\beta = 0.9$, $\gamma = 5$	11.21	3.28	3.07	2.56
$C = 5$, $\beta = 0.9$, $\gamma = 3$	10.39	4.30	4.23	3.89
$C = 5$, $\beta = 0.8$, $\gamma = 3$	11.91	4.45	4.71	4.20
$C = 5$, $\beta = 0.8$, $\gamma = 5$	13.92	4.47	4.59	4.15
$C = 1$, $\beta = 0.9$, $\gamma = 3$	8.79	5.37	4.97	4.36
$C = 2.5$, $\beta = 0.9$, $\gamma = 3$	7.07	4.78	4.36	3.57
$C = 2.5$, $\beta = 0.8$, $\gamma = 3$	10.5	6.54	5.35	4.35

have been discovered in many other asset markets including the exchange market. In particular, in the latter the returns have a kurtosis typically exceeding 3 and a measure of fat tails (Hill index) ranging between 2 and 5 (see Koedijk et al., 1992). However, it has also been detected that the kurtosis is reduced under time aggregation (Lux, 1998). This phenomenon has been observed for most exchange rates. We checked whether this is also the case with the simulated exchange rate changes in our model.[20]

The model was simulated using normally distributed random disturbances (with mean = 0 and standard deviation = 1). We computed the kurtosis and the Hill index of the simulated exchange rate returns. We computed the Hill index for 5 different samples of 2000 observations. In addition, we considered three different cut-off points of the tails (2.5%, 5%, 10%). We show the results of the kurtosis and of the Hill index in Table 4. We find that for a broad range of parameter values the kurtosis exceeds 3 and the Hill index indicates the presence of fat tails. Finally we check if the kurtosis of our simulated exchange rate returns declines under time aggregation. In order to do so, we chose different time aggregation periods and we computed the kurtosis of the time-aggregated exchange rate returns. We found that the kurtosis declines under time aggregation. In Table 5 we show the results for different parameter values (including low values for the transactions costs C).

These results suggest that the non-linear dynamics of the model transforms normally distributed noise in the exchange rate into exchange rate movements with tails that are significantly fatter than the normal distribution and with more density around the mean. Thus our model mimics an important empirical regularity, i.e. that exchange rate movements are characterised by tranquil periods (occurring most of the time) and turbulent periods (occurring infrequently).

[20]It should be noted that models similar to ours have been applied in the stock markets. These models have been able to replicate fat tails and excess kurtosis observed in these markets. See Hommes (2001).

Table 5
Kurtosis under time-aggregation

Parameter values	1 Period returns	10 Periods returns	25 Periods returns	50 Periods returns
$C = 5,\ \beta = 0.9,\ \gamma = 5$	11.21	4.75	3.79	4.73
$C = 5,\ \beta = 0.9,\ \gamma = 3$	10.39	9.93	3.19	2.28
$C = 5,\ \beta = 0.8,\ \gamma = 3$	11.91	17.69	3.01	2.97
$C = 5,\ \beta = 0.8,\ \gamma = 5$	13.92	14.65	3.86	2.85
$C = 5,\ \beta = 1,\ \gamma = 1$	9.32	9.01	9.07	2.56
$C = 1,\ \beta = 0.9,\ \gamma = 3$	8.79	8.56	4.01	3.78
$C = 2.5,\ \beta = 0.9,\ \gamma = 3$	7.07	6.43	3.22	2.45
$C = 2.5,\ \beta = 0.8,\ \gamma = 3$	10.5	13.2	2.92	2.67

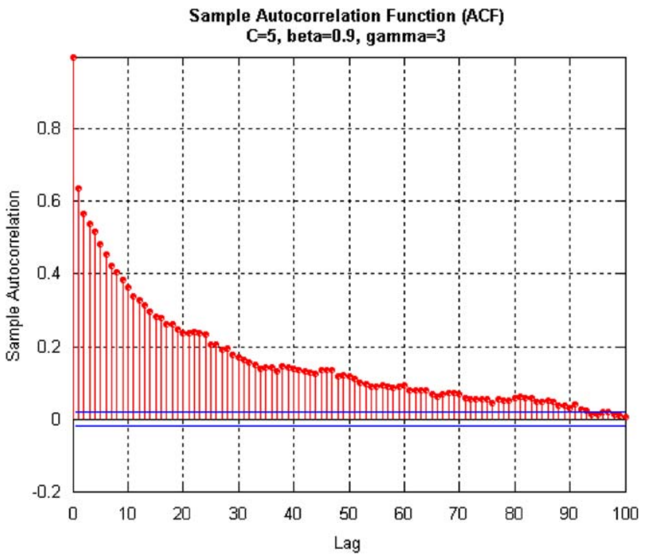

Fig. 12.

8.3. Volatility clustering

The last empirical regularity we investigate concerns the clustering of volatility. It has been widely observed that the exchange rate returns show a GARCH structure, i.e. there is time dependency in the volatility of the exchange rate returns (see Lux and Marchesi, 2000; Hommes, 2001; Kirman and Theyssière, 2002). In order to check if our model is capable of reproducing this statistical property we tested for GARCH structures in the simulated exchange rate returns. We first computed the autocorrelation function (ACF) of the absolute returns of the simulated exchange rate returns for a broad range of parameter values. In Fig. 12 we show the ACF for a particular set of parameters. At first glance, Fig. 12 suggests that the ACF dies out

Table 6
Serial correlation tests residuals ARMA(2,1)

	Model I		Model II	
	F-value	p-value	F-value	p-value
Breusch–Godfrey LM test	14.8	0.000	10.5	0.005
ARCH test	68.03	0.000	71.7	0.000

Model I = C = 5, $\beta = 0.9$, $\gamma = 3$; Model II = C = 2.5, $\beta = 0.9$, $\gamma = 3$.

Table 7
GARCH model

	I	II	III	IV
a	0.003 (2.6)	−0.003 (−2.5)	0.004 (3.5)	0.002 (1.2)
b	0.003 (14.5)	0.002 (14.1)	0.003 (15.8)	0.003 (12.3)
α	0.35 (19.1)	0.33 (31.8)	0.26 (22.7)	0.34 (15.6)
δ_1	0.44 (8.5)	0.49 (18.5)	0.41 (8.2)	0.37 (6.3)
δ_2	0.14 (3.7)	0.18 (5.6)	0.18 (4.6)	0.18 (4.1)

I: $C = 5$, $\beta = 0.9$, $\gamma = 3$; II: $C = 5$, $\beta = 0.9$, $\gamma = 5$; III: $C = 5$, $\beta = 0.8$, $\gamma = 3$; IV: $C = 2.5$, $\beta = 0.9$, $\gamma = 3$. Numbers in parentheses are t-statistics.

slowly, i.e. that the volatility in the exchange rate returns has a long memory. In order to confirm whether this visual impression is correct, we proceed as follows.[21] First we compute the ACF for the raw returns (see Fig. 14 in Appendix A). The wave-shape suggests an ARMA process for the returns. Thus, we estimate an ARMA model and we find that an ARMA(2,1) performs best. (The results of estimating this model on the returns are also shown in Appendix A.) We then tested for serial correlation in the residuals. These tests are shown in Table 6 for two different sets of parameters. We conclude that we should reject serial correlation in the error term.

The next step consisted in testing for GARCH effects in the exchange rate returns. In order to do so, we chose a GARCH (2,1) specification:[22]

$$\Delta s_t = a + \varepsilon_t,$$
$$\sigma_t^2 = b + \alpha\varepsilon_{t-1}^2 + \delta_1\sigma_{t-1}^2 + \delta_2\sigma_{t-2}^2,$$

[21]We also computed the rate of decay of the autocorrelation function by estimating an equation

$$\rho(k) = k^{d-1}$$

where $\rho(k)$ is the autocorrelation coefficient at lag k. We applied this to the absolute returns. We found that $d = 0.48$ and significantly different from zero, suggesting a significant departure from exponential decay. (Note that $d = 0$ for exponential decay, and $d = 1$ for uniform distribution of the autocorrelations.)

[22]We also estimated a GARCH (1,1). The results are very similar.

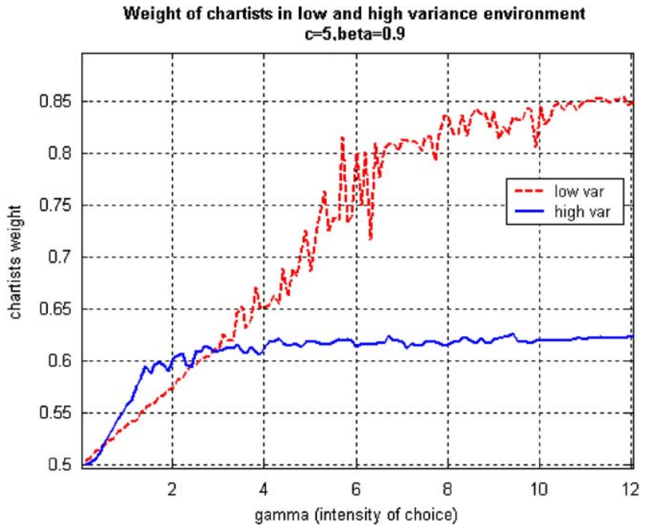

Fig. 13. *Note*: In high variance regime the variance of the shocks in the fundamentals is 10 times the variance in the fundamentals of the low variance regime.

where ε_t is the error term, a is a constant and σ_t^2 is the conditional variance of the returns. We estimated this model using the simulated exchange rate returns. We present the results in Table 7 for different values of the extrapolation parameter β.

We observe that the GARCH coefficients, α, δ_1 and δ_2, are significantly different from zero implying that there is volatility clustering in the exchange rate returns. In addition, we find that for values of β close to 0.9 the sum of α, δ_1 and δ_2, which is a measure of the degree of inertia of the volatility, is close to one. This implies that the effect of volatility shocks dies out slowly. Thus, our model is capable of reproducing a widely observed phenomenon of clustering and persistence in volatility.

9. Is chartism evolutionary stable?

An important issue is whether chartism survives in our model. Put differently, we ask the question under which conditions chartism is profitable such that it does not disappear. It should be noted that there is a broad literature that shows that technical analysis is used widely, also by large players (see Wei and Kim, 1997).

We investigate this issue by analysing how chartism evolves under different conditions. In Fig. 13 we show the average chartists' weight for increasing values of the intensity of choice parameter γ in two different environments concerning the variance of the shocks in fundamentals, a low and a high variance regime.[23] We obtained the chartists weights by simulating the model over 10 000 periods and

[23]We obtain qualitatively the same results when we keep the variance of the shocks constant while varying the transactions costs C. What matters is the size of the shocks relative to transactions costs.

computing the average weight over the last 5000 periods. Our first finding is that chartism does not disappear, i.e. in all simulations for many different parameters configurations we find that the weight attached to chartists never goes to zero. Second, for a wide range of parameter values we find that the chartists' weight fluctuates around a market share, which exceeds 50%. For high values of γ the chartists' weight approaches 90%. Third, when the shocks are large relative to the transactions cost band, the weight of chartists is considerably lower. This is in line with our previous results. When the fundamental shocks are large relative to transactions costs, the exchange rate is often driven outside the transaction cost band. As a result the goods market dynamics will often be operative, making fundamentalist forecasting relatively profitable. This reduces the scope for chartism.

These results are consistent with the empirical evidence of the importance of chartism in foreign exchange market (Taylor and Allen, 1992). And they also suggest that chartism is evolutionary stable.

10. Conclusion

In this paper we developed an exchange rate model, which has the following features. First, it introduces a non-linearity in the dynamics of the foreign exchange market that finds its origin in the existence of transaction costs in the goods markets. Second it allows for heterogeneity of the agents' beliefs. In particular, it assumes that agents use different forecasting rules, and that they switch to the most profitable one after evaluating their relative profitability. The model does not assume rational expectations. The problem of rational expectations models with heterogeneous agents is that it creates an 'infinite regress' problem, thereby imposing an unreasonable informational burden on individual agents. Therefore, we find it more useful to assume that in a highly complex world, agents use simple forecasting rules and evaluate the 'fitness' of these rules ex post.

The model generates a multitude of fixed-point attractors depending on the initial conditions. By adding exogenous noise the model produces a dynamics that resembles a chaotic one, although the deterministic part of the model is not chaotic. This feature has interesting implications. First, there is sensitivity to initial conditions, which implies that a small disturbance can drive the exchange rate on a different path. Second, the effect of a permanent shock in the fundamental exchange rate is largely unpredictable, i.e. one cannot forecast how the shock will affect the exchange rate in any particular point of time, but one can predict the *average* effect. It also implies that the exact timing of the shock matters. History matters.

The empirical relevance of the model is a measure of its quality. Therefore, we analysed to what extent our model is capable of reproducing the exchange rate puzzles that we observe in reality. The first puzzle we analysed is the 'disconnect puzzle'. This puzzle relates to the fact that the exchange rate movements are disconnected, most of the time, from the movements of the underlying fundamental variables. In our model 'disconnection' is a natural outcome when the variance of the underlying fundamentals is small compared to the size of the transactions costs in the goods market. We argued that this is the regime in which the currencies of the main industrialised countries find themselves in. In contrast, when the size of the fundamental shocks is large relative to the size of the transactions costs, the exchange rate is less disconnected from the underlying fundamentals.

Second, fat tails and excess kurtosis, which have been detected in the exchange rate returns, are generated by our model. In other words, our model generates a dynamics of the exchange rate with *intermittency* of high and low turbulence periods.

A third empirical regularity concerns the volatility clustering and persistence of exchange rate returns. We found GARCH effects in the simulated exchange rate returns that come close to the observed GARCH effects in the real life exchange rate returns.

A fourth empirical regularity is the continuing existence of chartists. This cannot easily be rationalised in the efficient market rational expectations model. In our model where chartists and fundamentalists continuously switch to the most profitable forecasting rule, chartists tend to dominate the market.

Some implications of these findings are the following. The exchange rates of the major currencies are subject to relatively small shocks in the underlying fundamentals (e.g. inflation differentials are almost zero). Compared to these shocks the transactions costs in the goods markets can be said to be relatively large (see Obstfeld and Rogoff, 2000, on this), i.e. a large part of goods and services are non-traded (or difficult to trade) because the cost of shipping them across borders is quite high. Thus, the regime confronted by the exchange rates of the major industrialised countries comes close to the regime we have identified to be the one where exchange rates are disconnected from fundamentals, and where excess volatility and speculative noise is produced by chartists' activity. Put differently, the movements of the exchange rates of the industrialised countries are likely to be clouded by a non-linear speculative dynamics that makes it difficult if not impossible to explain this or that movement of these exchange rates.

Acknowledgements

We are grateful to Volker Böhm, Yin-Wong Cheung, Hans Dewachter, Roberto Dieci, Marc Flandreau, Cars Hommes, Philip Lane, Thomas Lux, Richard Lyons, Michael Moore, Ronald McDonald, Assaf Razin, Piet Sercu, Peter Sinclair, Jan Tuinstra and Peter Westaway for comments and suggestions. The comments of two anonymous referees are also gratefully acknowledged.

Appendix A. Diagnosis of GARCH structure in the simulated exchange rate returns

In this Appendix we present the ACF of the simulated exchange rate returns (Fig. 14) and the ARMA(2,1) estimation on these returns (Table 8).

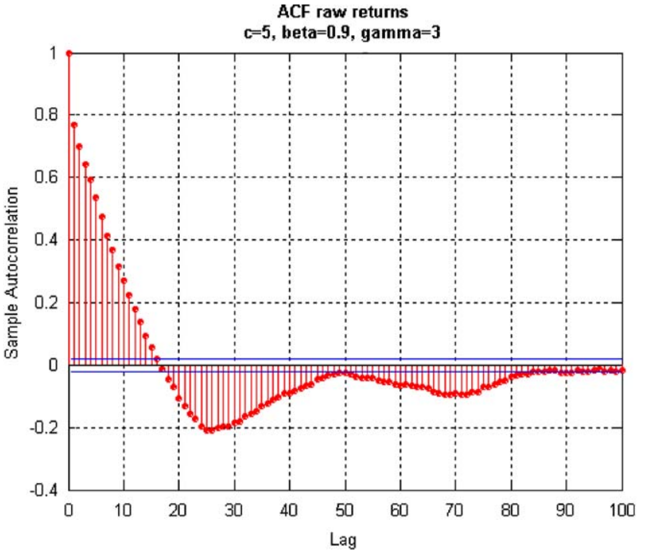

Fig. 14. ACF of simulated exchange rate returns.

Table 8
Estimation of ARMA(2,1) on simulated returns

Variables	Coefficient	*t*-statistic
Constant	−0.001	−0.14
AR(1)	0.856	28.74
AR(2)	0.044	1.83
MA(1)	−0.362	−12.71
R^2	*0.60*	

References

Bacchetta, P., van Wincoop, E., 2003. Can information heterogeneity explain the exchange rate determination puzzle? *NBER* working paper 9498.

Baxter, M., Stockman, A., 1989. Business cycles and the exchange rate regime. Some international evidence. Journal of Monetary Economics 23, 377–400.

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

Brock, W., Hommes, C., 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. Journal of Economic Dynamics and Control 22, 1235–1274.

Cheung, Y., Chinn, M., Marsh, I., 1999. How do UK based foreign exchange dealers think their markets operate? CEPR, Discussion Paper, no. 2230, London.

Cheung, Y., Lai, K., Bergman, M., 2001. Dissecting the PPP puzzle: the unconventional roles of nominal exchange rate and price adjustments. Paper presented at CES-Ifo Conference, Munich, 2002.

Chiarella, C., Dieci, R., Gardini, L., 2002. Speculative behaviour and complex asset price dynamics. Journal of Economic Behaviour and Organisation 49, 173–197.

De Boeck, J., 2000. The effect of macroeconomic 'news' on exchange rates, a structural VAR approach. University of Leuven, Mimeo.

De Grauwe, P., Grimaldi, M., 2001. Exchange rates, prices and money: a long run perspective. International Journal of Finance and Economics 6 (4), 289–314.

De Grauwe, P., Vansteenkiste, I., 2001. Exchange rates and fundamentals. A non-linear relationship? CESifo Working Paper, no. 577, October.

De Grauwe, P., Dewachter, H., Embrechts, M., 1993. Exchange Rate Theories. Chaotic Models of the Foreign Exchange Markets. Blackwell, Oxford.

de Vries, C., 2001. Fat tails and the history of the guilder. Tinbergen Magazine 4, 3–6.

De Long, J., Bradford, B., Shleifer, Summers, L., 1990. Noise trader risk in financial markets. Journal of Political Economy.

Dumas, B., 1992. Dynamic equilibrium and the real exchange rate in a spatially separated world. Review of Financial Studies 5 (2), 153–180.

Engel, C., 2000. Long run PPP may not hold after all. Journal of International Economics 57.

Engel, C., Rogers, J., 1995. How Wide is the Border. International Finance Discussion Paper, no. 498, Board of Governors of the Federal Reserve System, Washington, DC.

Faust, J., Rogers, J., Wright, J., 2003. Exchange rate forecasting: the errors we've really made. International Finance Discussion Papers 741, Federal Reserve Board of Governors, Washington, DC.

Flood, R., Rose, A., 1995. Fixing the exchange rate regime: a virtual quest for fundamentals. Journal of Monetary Economics 36, 3–37.

Frankel, J., Froot, K., 1986. The dollar as a speculative bubble: a tale of fundamentalists and chartists. NBER Working Paper, no. 1963.

Goodhart, C., 1989. News and the foreign exchange market. LSE Financial Markets Group Discussion paper, 71.

Goodhart, C., Figlioli, L., 1991. Every minute counts in the foreign exchange markets. Journal of International Money and Finance 10, 23–52.

Guillaume, D., 1996. Chaos, randomness and order in the foreign exchange markets. Ph.D. Thesis, K.U. Leuven, Leuven.

Haskel, J., Wolf, H., 2001. The law of one price. A case study. CESifo Working Paper, no. 428, March.

Hommes, C., 2001. Financial markets as complex adaptive evolutionary systems. Quantitative Finance 1, 149–167.

Johansen, S., 1991. Estimation and hypothesis testing of cointegration vectors in Gaussian vector autoregressive models. Econometrica 55, 1551–1580.

Kilian, L., Taylor, M., 2001. Why is it so difficult to beat the random walk forecast of exchange rates? University of Warwick, Mimeo, pp. 29.

Kirman, A., Theyssière, G., 2002. Microeconomic models for long memory in the volatility of financial time series. Studies in Nonlinear Dynamics and Econometrics 5 (4), 281–302.

Koedijk, K., Stork, P., de Vries, C., 1992. Differences between foreign exchange rate regimes: the view from tails. Journal of International Money and Finance 11, 462–473.

Lux, T., 1998. The socio-economic dynamics of speculative markets: interacting agents, chaos, and fat tails of return distributions. Journal of Economic Behaviour and Organisation 33.

Lux, T., Marchesi, M., 2000. Volatility clustering in financial markets: a microsimulation of interacting agents. International Journal of Theoretical and Applied Finance 3, 675–702.

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

Meese, R., Rogoff, 1983. Empirical exchange rate models of the seventies: do they fit out of sample?. Journal of International Economics 14, 3–24.

Michael, P., Nobay, R., Peel, A., 1997. Transaction costs and non-linear adjustment in real exchange rates: an empirical investigation. Journal of Political Economy 105 (4), 862–879.

Obstfeld, M., Rogoff, K., 1996. Foundations of International Macroeconomics. MIT Press, Cambridge, MA.

Obstfeld, M., Rogoff, K., 2000. The six major puzzles in international macroeconomics: is there a common cause? NBER Working Paper no. 7777, July.

Sarno, L., Taylor, M., 2002. The economics of exchange rates. Cambridge University Press, Cambridge 318pp.
Schittenkopf, C., Dorffner, G., Dockner, E., 2001. On nonlinear, stochastic dynamics in economics and financial time series. Studies in Nonlinear Dynamics and Econometrics 4 (3), 101–121.
Shleifer, A., 2000. Introduction to Behavioural Finance. Clarendon Press, Oxford.
Taylor, M., Allen, H., 1992. The use of technical analysis in the foreign exchange market. Journal of International Money and Finance 11, 304–314.
Taylor, M., Peel, D., Sarno, L., 2001. Non-linear mean reversion in real exchange rates: towards a solution to the purchasing power parity puzzles. CEPR discussion paper no. 2658.
Wei, S.-J., Kim, J., 1997. The big players in the foreign exchange market: do they trade on information or noise? NBER working paper, no. 6256.
Williamson, J., 1985. The Exchange Rate System, Policy Analyses in International Economics. 5, Institute for International Economics, Washington, DC.