The Exchange Rate in a Model with Heterogeneous Agents and Transactions Costs

Paul De Grauwe and Marianna Grimaldi*
University of Leuven

February 14, 2003

Abstract

In this paper we develop a model of the exchange rate. The existence of transactions costs introduces an important non-linearity. Agents have different beliefs about the future exchange rate. We show that this simple model creates great complexity in the market which is characterised by the fact that the exchange rate is disconnected from its fundamental most of the time. Periods of tranquility and turbulence alternate in unpredictable manner. Finally we show that this model mimicks most of the empirical puzzles uncovered in the literature.

*Corresponding address: Naamsestraat 69, 3000 Leuven, Belgium.

1 Introduction

This paper suggests a new framework for understanding the exchange rate dynamics. Recently, two new approaches have been introduced in the literature. One set of models has focused on the presence of non linearities. Transaction costs have been recognised as one of the most important sources of non linearities affecting the exchange rate dynamics. Another set of models has highlighted the role of the heterogeneity of agents. Although these two approaches have been growing mostly apart we believe that there is a need to combine these two insights for a better understanding of the exchange rate dynamics.

Using a mean-variance optimisation framework we develop a simple non linear exchange rate model with heterogeneous agents and assuming transactions costs in the goods markets. It has become customary in exchange rate models with heterogenous agents to assume that some agents are perfectly informed (have rational expectations) and others use imperfect information (noise traders). We depart from this setup. It seems to us that it is not reasonable to make such an asymmetric assumption because it leads immediately to the unresolved question of why some agents have perfect information, and others lack this. Instead it seems more reasonable to take the view that all agents are imperfectly informed. Therefore we will introduce the assumption that agents use a forecasting rule based on incomplete information. However, these agents continuously check the profitability of their forecasting rule and they are willing to switch to another rule if this turns to be more profitable (in a risk adjusted way). Thus a selection process based on fitness of the forecasting rules is imposed[1].

We show that our model is capable of replicating the empirical puzzles and anomalies uncovered in the last decade by the empirical exchange rate literature.

2 The model

In this section we develop a simple non-linear model of the exchange rate . We assume agents of different types i depending on their beliefs about the future exchange rate. Each agent can invest in two assets, a domestic and a foreign one. The agents' utility function can be represented by the following equation:

$$U(W_{t+1}^i) = E_t(W_{t+1}^i) - \frac{1}{2}\mu V^i(W_{t+1}^i) \tag{1}$$

where W_{t+1}^i is the wealth of agent of type i at time t+1, E_t.is the expectation operator, μ is the coefficient of risk aversion and $V^i(W_{t+1}^i)$ represents the conditional variance of wealth of agent i. The wealth is specified as follows:

$$W_{t+1}^i = (1 + r^*)\, s_{t+1} d_t^i + 1 + r\left(W_t^i - s_t d_t^i\right) \tag{2}$$

[1]See Brock and Hommes(1998) for such an approach in a stock market model. See also Kurz(1994) and Kurz and Motolese(2000).

where r and r^* are respectively the domestic and the foreign interest rates, s_{t+1} is the exchange rate at time t+1, $d_{i,t}$ represents the holdings of the foreign assets by agent of type i at time t. Thus, the first term on the right-hand side of 2 represents the value of the foreign portfolio in domestic currency at time t+1 while the second term represents the value of the domestic portfolio at time t+1.

Substituiting equation 2 in 1 and maximising the utility with respect to $d_{i,t}$ allows us to derive the optimal holding of foreign assets by agents of type i :

$$d_{i,t} = \frac{(1+r^*)\, E_t^i\,(s_{t+1}) - (1+r)\, s_t}{\mu \sigma_{i,t}^2} \tag{3}$$

The market demand for foreign assets at time t is the sum of the individual demands, i.e.:

$$\sum_{i=1}^{N} n_{i,t} d_{i,t} = D_t \tag{4}$$

where $n_{i,t}$ is the number of agents of type i.

Market equilibrium implies that the market demand is equal to the market supply $X_t$which we assume to be exogenous[2]. Thus,

$$X_t = D_t \tag{5}$$

Substituting the optimal holdings into the market demand and then into the market equilibrium equation and solving for the exchange rate s_t yields the equilibrium exchange rate:

$$s_t = \left(\frac{1+r^*}{1+r}\right) \frac{1}{\sum\limits_{i=1}^{N} \frac{n_{i,t}}{\sigma_{i,t}^2}} \left[\sum_{i=1}^{N} n_{i,t} \frac{E_t^i\,(s_{t+1})}{\sigma_{i,t}^2} - \mu \frac{X_t}{1+r}\right] \tag{6}$$

In order to model the expectations formation we assume that there are two types of agents: chartists and fundamentalists. As a result equation 6 specialises to :

$$s_t = \left(\frac{1+r^*}{1+r}\right) \frac{1}{\left(\frac{n_{f,t}}{\sigma_{f,t}^2} + \frac{n_{c,t}}{\sigma_{c,t}^2}\right)} \left[n_{f,t} \frac{E_t^f\,(s_{t+1})}{\sigma_{f,t}^2} + n_{c,t} \frac{E_t^c\,(s_{t+1})}{\sigma_{c,t}^2} - \mu \frac{X_t}{1+r}\right] \tag{7}$$

[2]The market supply is determined by the net current account and by the sales or purchases of foreign exchange of the central bank. We assume both to be exogenous. In an extension of this paper we intend to endogenise the market supply.

Thus the exchange rate is determined by the expectations of fundamentalists and chartists about the future exchange rate. These forecasts are weighted by their respective variances.

We now specify how fundamentalists and chartists.form their expectations of the future exchange rate. In a second step we will specify how they take into account the risk as measured by the variances.

The fundamentalists base their forecast on a comparison between the market and the fundamental exchange rate, i.e. they forecast the market rate to return to the fundamental rate in the future. In this sense they use a negative feedback rule that introduces a mean reverting dynamics in the exchange rate . The speed with which the market exchange rate returns to the fundamental is assumed to be determined by the speed of adjustment in the goods market. Thus, the forecasting rule for the fundamentalists is :

$$E_t^f\left(\Delta s_{t+1}\right) = -\psi\left(s_t - s_t^*\right) \tag{8}$$

where s_t^* is the fundamental exchange rate at time t , which is assumed to follow a random walk.

The chartists forecast the future exchange rate by extrapolating past exchange rate movements. Their forecasting rule can be specified as :

$$E_t^c\left(\Delta s_{t+1}\right) = \beta \sum_{i=0}^{T} \alpha_i \Delta s_{t-i} \tag{9}$$

Thus, the chartists compute a moving avrage 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 in contrast to the fundamentalists they do not take into account information concerning the fundamental exchange rate. In this sense they can be considered to be pure noise traders .In a way this chartist rule can also be seen as reflecting herding behaviour, i.e. chartists do not take fundamental information into account because they feel to uncertain about their meaning, but they closely watch the movements of the exchange rate as a way to detect "market sentiments". If the latter are positive, they buy; if they are negative, they sell.

Our choice to give a prominent role to chartists' rules of forecasting is based on empirical evidence. The evidence that Chartism is used widely to make forecasts is overwhelming (see Cheung and Chinn(1989), Taylor and Allen(1992)). 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. We return to this issue when we let the number of chartists be determined by the profitability of the chartists' forecasting rule.

We now analyse how fundamentalists and chartists evaluate the risk. The latter is measured by the variance terms in equation 7, which we define as the

weighted average of the squared (one period ahead) forecasting errors made by chartists and fundamentalists, respectively. Thus,

$$\sigma_{i,t} = \sum_{k=1}^{\infty} \gamma_k \left[E^i_{t-k} \left(s_{t-k+1} \right) - s_{t-k+1} \right]^2 \tag{10}$$

where γ_k are geometrically declining weights.

However fundamentalists and chartists perceive the risk in a different way. In particular the fundamentalists are assumed to take into account the deviation of the exchange rate from the fundamental in addition to the forecasting error. We will call the deviation of the market exchange rate from its fundamental, the misalignment. Thus the fundamentalists'risk term can be written as:

$$\sigma_{f,t} = \frac{\sum_{k=1}^{T} \gamma_k \left[E^f_{t-k} \left(s_{t-k+1} \right) - s_{t-k+1} \right]^2}{\left(s_{t-1} - s^*_{t-1} \right)^2} \tag{11}$$

where $\left(s_{t-1} - s^*_{t-1} \right)$ is the misalignment.

The logic behind this specification is that the fundamentalists consider the fundamental exchange rate as a benchmark. The larger is the misalignment the less the fundamentalists will attach importance to the short term volatility as measured by the one-period ahead forecasting error. Put differently, the larger is the deviation of the exchange rate from its fundamental, the smaller is the risk of using a fundamentalist forecasting rule, given the past (short-term) forecast errors made by the fundamentalsits. In contrast the chartists do not take into account the misalignment.

We now specify the dynamics that governs the number of chartists and fundamentalists, namely n_{ct} and n_{ft}. In order to do so, we describe how the number of chartists and fundamentalists changes from period t-1 to period t:

$$n^c_t = n^c_{t-1} + n^f_{t-1} p^{fc}_t - n^c_{t-1} p^{cf}_t \tag{12}$$

$$n_{f,t} = n_{f,t-1} + n_{c,t-1} p^{cf}_t - n_{f,t-1} p^{fc}_t \tag{13}$$

where $n_{c,t}$ and $n_{f,t}$ are the number of chartists and fundamentalists in period t; p^{cf}_t represents the probability of a chartist to become a fundamentalist in period t, and p^{fc}_t is the probability of a fundamentalist to become a chartist in period t.

These probabilities are assumed to be a function of the relative utilities of chartism and fundamentalism. These utilities depend on the profitability of the

forecasting rules and the risk associated with their use. The probabilities are specified as follows:

$$p_t^{fc} = \varsigma \frac{exp\left(U_{c,t-1}\right)}{exp\left(U_{c,t-1}\right) + exp\left(U_{f,t-1}\right)} \tag{14}$$

$$p_t^{cf} = \varsigma \frac{exp\left(U_{f,t-1}\right)}{exp\left(U_{c,t-1}\right) + exp\left(U_{f,t-1}\right)} \tag{15}$$

where $U_{c,t-1}$ and $U_{f,t-1}$ are the utility functions of the chartists and fundamentalists in period t-1. These probabilities of switching can be interpreted as follows. When the utility associated with using chartist information increases relative to the utility of using fundamentalist information the agents tend to switch from fundamentalism to chartism and vice versa. Taking into account the specification of the utility function (equation 1), we can rewrite equations 14 and 15 as follows:

$$p_t^{fc} = \varsigma \frac{exp\left(\pi_{c,t-1} - \mu\sigma_{c,t-1}^2\right)}{exp\left(\pi_{c,t-1} - \mu\sigma_{c,t-1}^2\right) + exp\left(\pi_{f,t-1} - \mu\sigma_{f,t-1}^2\right)} \tag{16}$$

$$p_t^{cf} = \varsigma \frac{exp\left(\pi_{f,t-1} - \mu\sigma_{f,t-1}^2\right)}{exp\left(\pi_{c,t-1} - \mu\sigma_{c,t-1}^2\right) + exp\left(\pi_{f,t-1} - \mu\sigma_{f,t-1}^2\right)} \tag{17}$$

Thus, the probability of switching depends on the relative profits adjusted for the risk. When the risk adjusted profits of chartist rules increase relative to the risk adjusted profits of fundamentalist rules, the probability of switching from fundamentalism to chartism increases, and vice versa. Agents make a profit when they correctly forecast the direction of the exchange rate movements. They make a loss if they wrongly predict the direction of its movements. The profit (the loss) they make equals the one-period return of investing \$1. The parameter ς measures the sensitivity of the probability of switching with respect to the relative profits adjusted for the risk . Note also that $0 < \varsigma < 1$.

3 The model with transactions costs

There is an increasing body of theoretical literature stressing the importance of transactions costs as a source of non-linearity in the determination of the exchange rate (Dumas(1992), Sercu, Uppal and Van Hulle(1995), Obstfeld and Rogoff(2000)). The importance of transaction costs has also been confirmed empirically (Taylor, Peel, and Sarno(2001), Kilian and Taylor(2001)). Therefore, we will develop a version of the previous model in which the transaction costs play a role.

We take the view that if transaction costs exist, the fundamentalists take this information into account. Therefore, if the exchange rate is within the transaction costs band the fundamentalists behave differently than if the exchange rate moves outside the transaction costs band.

Consider the first case, when the exchange rate deviations from the fundamental value are smaller than the transaction costs. In this case the fundamentalists know that arbitrage in the goods market does not apply. As a result, they expect the changes in the exchange rate to follow a white noise process. The best they can do is to forecast no change. More formally,

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

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 equation 8. More formally,

$$\text{when} \mid s_t - s_t^* \mid > C \quad \text{holds,} \qquad \text{then equation 8 applies.}$$

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

4 Solution of the model

In this section we investigate the properties of the solution of the model. We first study its deterministic solution. This will allow us to analyse the characteristics of the solution that are 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. 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.. We illustrate this feature in figure 1 where we plot 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.Figure 1 shows the fixed point solutions when the speed of adjustment in the goods market is 0.2. Note the complex pattern of

Figure 1:

these fixed point solutions, with many discontinuities. This has the implication that a small change in the initial condition can have a large effect on the solution This feature lies at the heart of some of the results that are obtained with this model relating to the unpredictability of the effect of shocks in exogenous variables. We return to this phenomenon in section 6.

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' 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.

It is also interesting to point out the existence of critical points. These are often obtained close to the limit of the transactions cost band. A continuous line is then broken into a cluster of points. These have a fractal nature, in that they

have an infinite self-replicating characterisitics. We explore this characteristic further in appendix 1.

5 Sensitivity analysis

In this section we perform a sensitivity analysis. We do this by showing diagrams that relate the solutions to different values of important parameters of the model. We concentrate on the extrapolation parameter used by the chartists, as this turns out to be a crucial parameter determining the nature of the solutions[3]. Figures 2 and 3show examples of such diagrams. They are analogous to so-called bifurcation diagrams. On the vertical axis we set out different values of the extrapolation parameter. On the vertical axis we show the solutions for the exchange rate. This is the exchange rate obtained after 1000 periods, given an initial shock. The two diagrams were constructed for two different initial shocks. We observe the following. For values of $\beta \leq 0.8$ we obtain a continuous line. This means that there is a unique fixed point solution for each value of β.When β is approximately 0.8 we have a critical point where the continuous line breaks into a fractal pattern. We stil have fixed point solutions, though. The continuous line, however, is fractured in such a way that for each value of β only one solution is found. We show a blow-up of the region around this critical point in appendix 1. When β reaches a value of approximately 0.9, we enter the chaotic region. This is characterised by infinitely many solutions for each value of β.These points correspond to strange attractors within which the exchange rate then travels. Note that we do not obtain bifurcations in this model, like the Hoppf bifurcation. The transition to chaos is abrupt. We do find that increasing the value of β can lead us in and out of the chaotic region as is shown in Figure 4.

We have produced two diagrams corresponding to two different initial conditions (Figures 2 and 3). We observe that the initial conditions determine the path that leads the model into chaos. For each different initial condition there is a different path. In addition, it appears that the initial conditions also influence the dynamics in a more fundamental way. This can be seen from a comparison of Figures 2 and 3. We observe that the different initial condition leads the exchange rate to switch out of chaos and back into it when β is close to the borderline between fixed point solutions and chaos.

We investigated this feature further in Figure 4 where we set $\beta = 0.9$ and vary the initial conditions. We now observe that the different initial conditions lead to switches in and out of the chaotic region. These features illustrate the great complexity in the exchange rate dynamics. As we will be analysed in greater detail later, this complexity has many different implications. It acts as a veil obscuring the transmission of exogenous shocks (e.g. shocks in the fundamental exchange rate) into the market exchange rate. It has the potential of producing regime switches triggered by small disturbances. Finally, this complexity greatly

[3]In appendix we do a similar analysis varying other parameters.

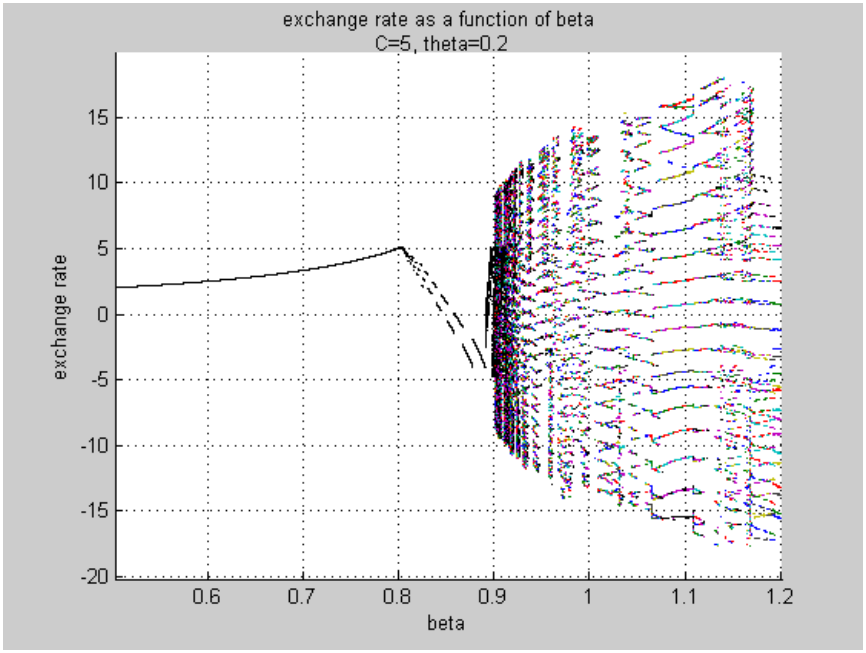


Figure 2:

complicates the making of standard statistical inferences from the distribution of the exchange rate changes.

We also investigated the importance of transaction costs. In order to do so, we set the transactions cost C=0, and produced similar bifurcation diagrams as in the previous paragraphs. We show the results in figure 5. A comparison of figure 5 with figures 2 and 3 reveals similarities and differences. The similarities are that when β reaches a critical value (close to 1) the fixed point solutions switch into chaotic solutions. In addition there is a similar dependence on initial conditions. This is made clear from a comparison of figure 6 with figure 1.

The differences are that in the absence of transactions costs the critical value of β that turns fixed point solutions into chaos is somewhat larger. In addition, the variability of the exchange rate is substantially smaller in the absence of transactions costs. This is also made clear from figure 7 where we produce a similar bifurcation diagram as in figures 2 and 3. We now allow the transactions costs C to vary (horizontal axis) while we fix $\beta = 1$. We observe that when the transactions costs increase the range of values we obtain for the exchange rate increases.

It is also important to analyse the dynamics of the weights of chartists and fundamentalists. We find that when the exchange rate converges towards a fixed-point attractor the weights of chartists and fundamentalists converge to 0.5. This can be explained as follows. When the exchange rates reach a

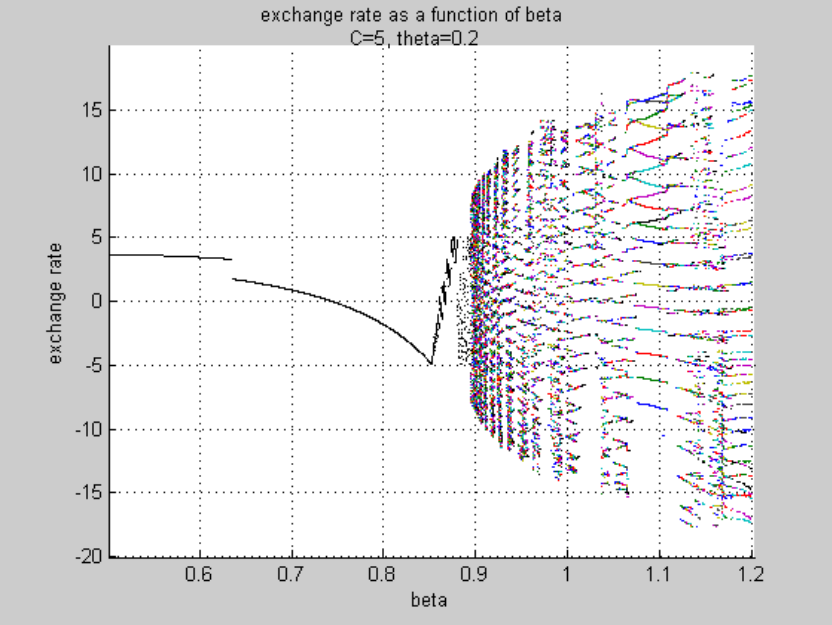


Figure 3:

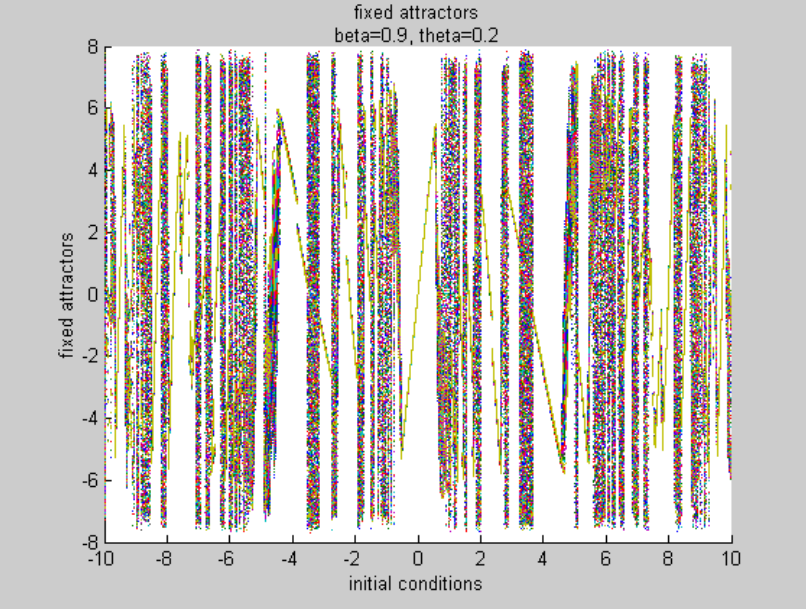


Figure 4:

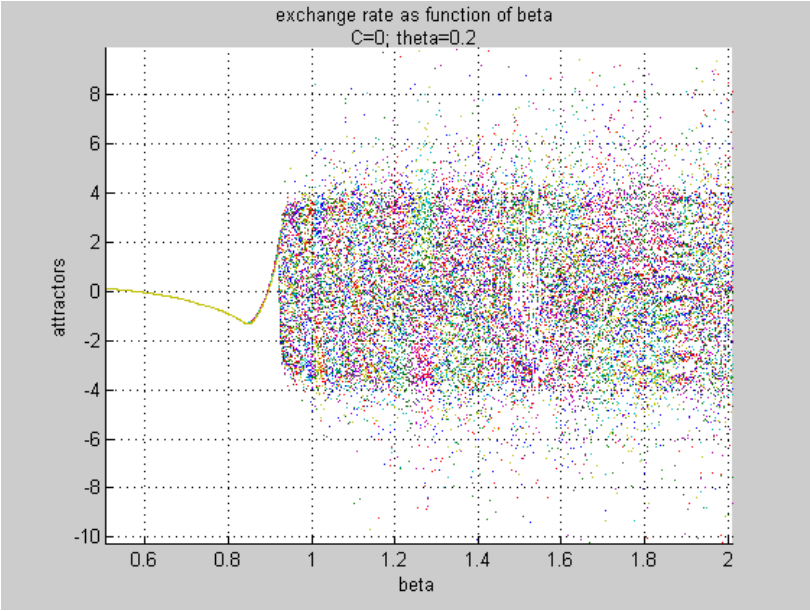


Figure 5:

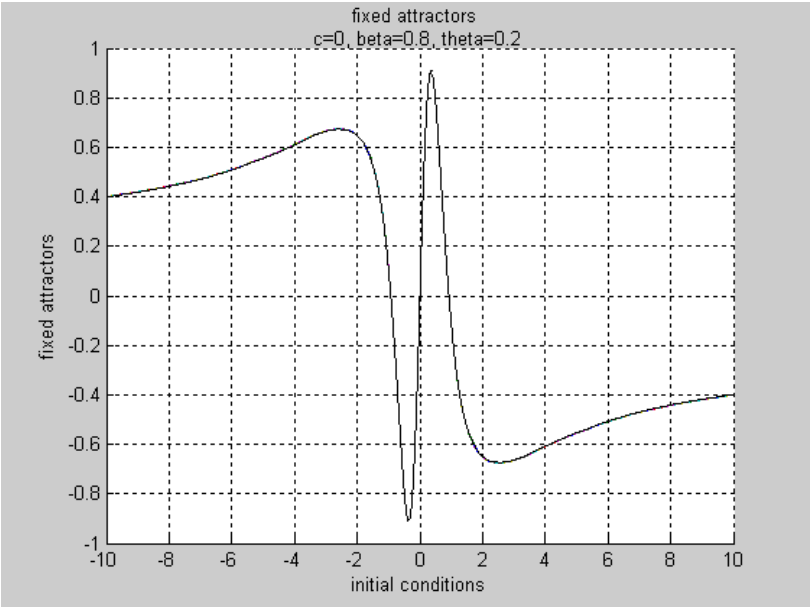


Figure 6:

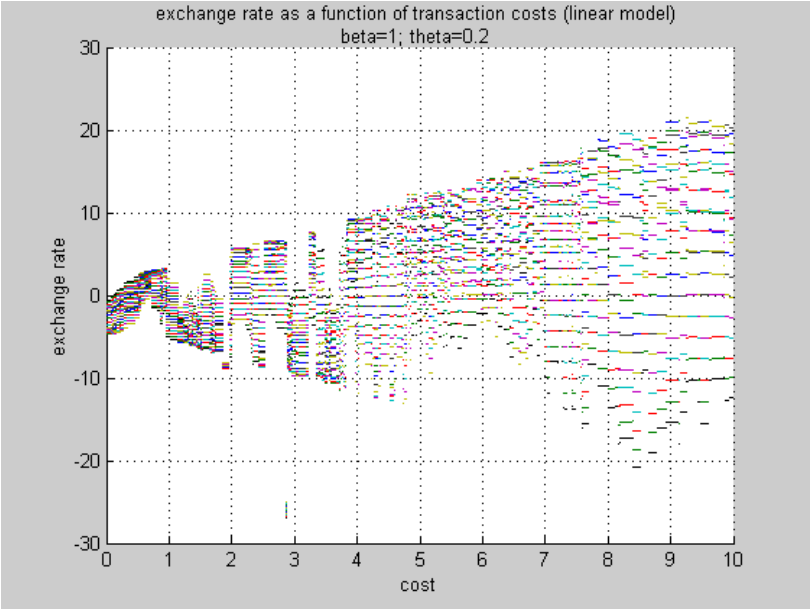


Figure 7:

fixed-point solution, chartists and fundamentalists expect no change anymore[4]. Therefore, they do not buy or sell, and thus make neither profits nor losses. As a result, the profit related selection rule of the model (see equations 11 and 12) assures that they will be equally represented in the market.

Things are very different when the exchange rate follows a chaotic pattern. In this case the chartists and fundamentalists weights will show cyclical movements. In general the chartists' weight will tend to fluctuate around a value larger than 0.5 provided the chartist extrapolation is not too high. We come back to this feature in section 8 where we analyze the dynamics of the weights of chartists and fundamentalists in more detail.

The empirical evidence in favour of deterministic chaos is not very strong. Sometimes deterministic chaos has been detected in the data, but most often no such dynamics has been found (D. Guillaume (1996), C. Schittenkopf, G. Dorffner and E. Dockner (2001)). This suggests that we should restrict the analysis of the model to parameter configurations that do not produce deterministic chaos. Typically this implies restrciting the extrapolation parameter used by

[4]Things are a little more complicated. When transactions costs are zero, we obain non-zero fixed points. Since the fundamental exchange rate was normalised at zero, this means that the market exchange rate is permanently different from the fundamental rate. Thus in the steady state the fundamentalists expect a constant appreciation (depreciation). However, because of the risk involved, fundamentalists do not act on this. The difference between the market rate and the fundamental rate can be considered to be a the risk premium. Our model generates risk premia whose value depends on the initial conditions.

chartists not to exceed 0.9 - 1 most of the time. The borerline between fixed point solutions and chaos remains interesting, though. As we have seen, even with an extrapolation parameter between 0.9 and 1 different initial conditions lead to swithches in and out of chaos, making it difficult to detect deterministic chaos from the data[5].

In the next section we restrict the analysis of the model for 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 complex dynamics that exhibits many of the features of chaotic dynamics despite the fact that the deterministic solutions of the model are fixed points.

6 The stochastic version of the model

We now introduce stochastic disturbances to the model. In our model these disturbances appear in two places. First they affect the fundamental, which is assumed to be a random walk. Second, as can be seen from equation (8), 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, β= 0.5 and θ= 0.2. (Similar results are obtained for a wide range of parameter values).

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 we simulated the model 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 figure 8.

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 figure 1,that 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 they are attracted by fixed- points that are far away from each other. The nice aspect of this is that we obtain a result that is typical for chaotic systems, however, without chaos being present in the deterministic part of the model. The combination of exogenous noise and a multiplicity of fixed-point attractors creates chaos-like dynamics.

A second feature of the model relates to the way shocks in the fundamental exchange rate are transmitted into 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

[5]Note that this switching in and out of chaos depending on the initial conditions tends to disappear for sufficiently high values of β.

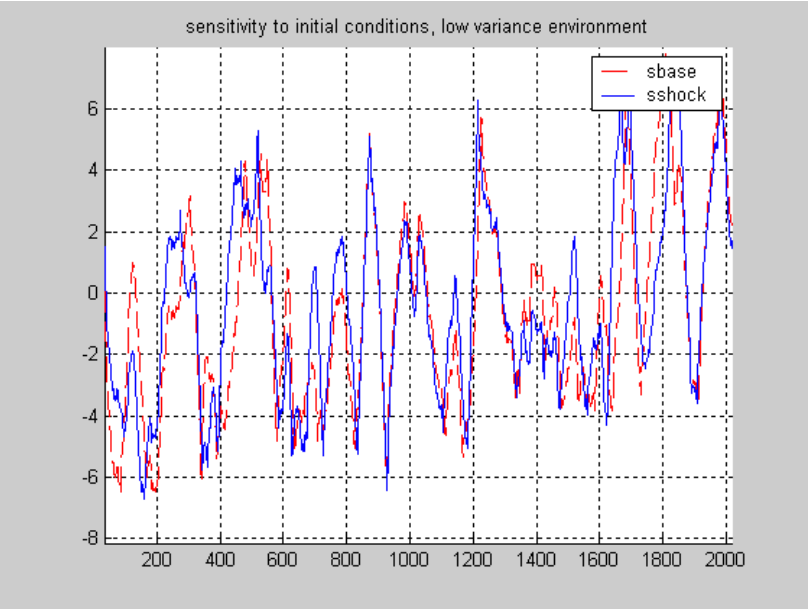


Figure 8:

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. In a linear model we would find that in the long run the exchange rate increases by 10. This is not the case in our model. We present the evidence in figure 9 where we show the effect of the same permanent shock of 10 in the fundamental rate on the exchange rate. The simulations are done assuming exactly the same stochastics in the scenario with as without the permanent shock in the fundamental exchange rate. Thus, there is no exogenous noise in the model that could blur the transmission process from the fundamental rate to the exchange rate.

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

We performed a sensitivity test of this result. In order to do so, we plot the effects of the same shock in the fundamental variable against the extrapolation parameter β. We show the result in figure 10. We also plot two lines representing

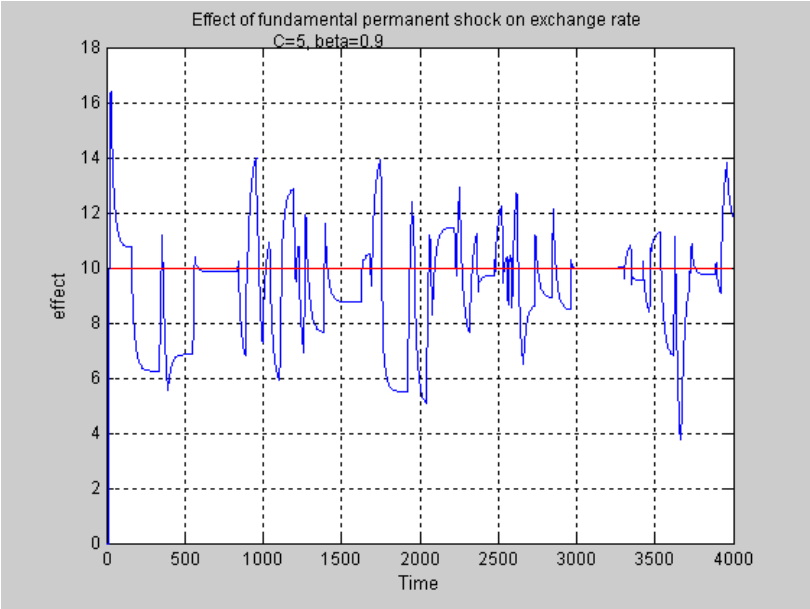


Figure 9:

plus and minus one standard deviation of the effects of the shock (the heavy lines) . We observe that as β increases the standard deviations increases, i.e. the precision with which one can forecast the effect of a given shock in the fundamental exchange rate declines.

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 figure 11.

We observe that the small difference in timing changes the whole 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.

Our results help to explain why in the real world it appears so 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

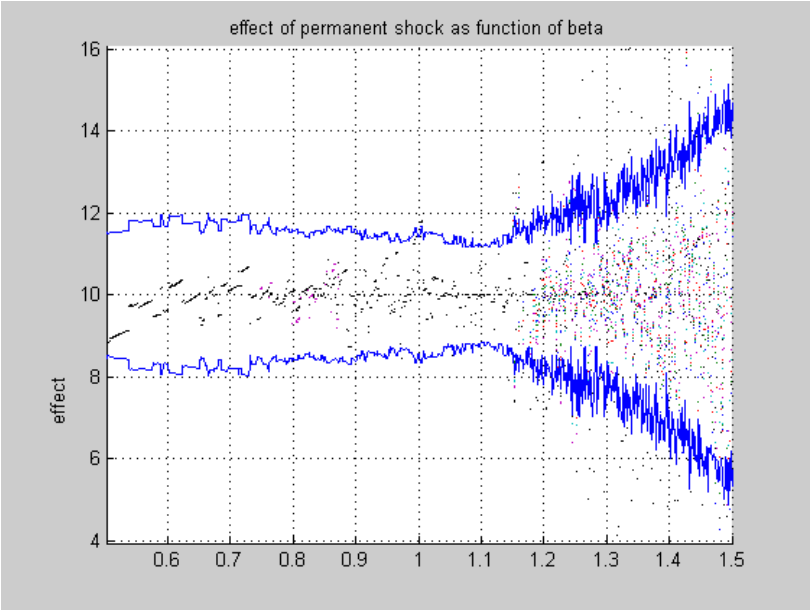


Figure 10:

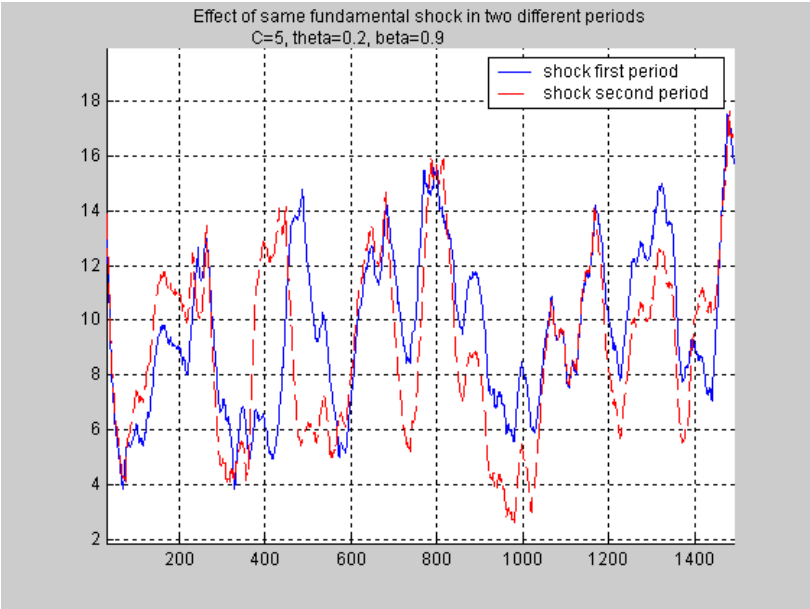


Figure 11:

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) matters 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.

Finally, it should be stressed that the uncertainty about the effect of a permanent shock in the fundamental only holds in a particular environment that is related to a low variance of the noise. In a later section we will analyse how different environments concerning the variance of shocks affect the results.

7 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 start with the 'disconnect puzzle'.

7.1 The disconnect puzzle

The first and foremost empirical puzzle has been called the "disconnect" puzzle (see Obstfeld and Rogoff(2000)), i.e. the exchange rate appears to be disconnected from its underlying fundamentals most of the time. 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. Goodhart (1989) and Goodhart and Figlioli (1991) found that most of the changes in the exchange rates occur when there is no observable news in the fundamental economic variables. This finding contradicted the theoretical models (based on the efficient market hypothesis), which imply that the exchange rate can only move when there is news in the fundamentals.

Our model is capable of mimicking this empirical regularity. In figure 12 we show the market exchange rate and the fundamental rate for a combination of parameters that does not produce deterministic chaos. (Our results hold equally well for a large set of parameter values including those that 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

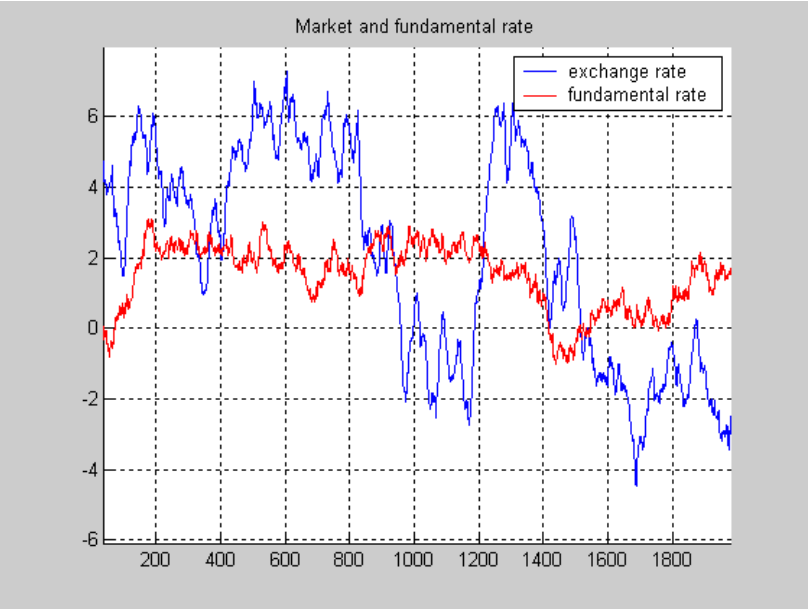


Figure 12:

fundamental using the same parameter values as in figure 12 for a sample of 8000 periods. We found that there is a cointegration relationship between the exchange rate and its fundamental . 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 a EC model in the following way:

$$\Delta s_t = \mu \left(s_{t-1} - \gamma s^*_{t-1} \right) + \sum_{i=1}^{n} \lambda_i \Delta s_{t-i} + \sum_{i=1}^{n} \phi_i \Delta s^*_{t-i} \tag{18}$$

The first term on the right hand side is the error correction term. The result of estimating this equation is presented in figure **??** where we have set n=4 .

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 0.5% 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 to the chartists' extrapolation behaviour. This phenomenon has been observed in reality. Chen et al.(2002) have recently discovered that most of the

Figure 13: Parameter estimates of VEC model (equation 19)

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
-0.005	1.06	0.18	0.15	0.09	0.05	0.04	0.03	0.02	0.01
-6. 4	*11.3*	*16.4*	*13.1*	*7.06*	*4.3*	*2.2*	*1.6*	*1.3*	*1.06*

slow mean reversion of the real exchange rate is due to slow adjustment of the nominal exchange rate and not of the goods prices.

From figure **??**, 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 exchange rate . 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))[6] .

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. Our model closely mimics these empirical regularities.

7.2 The " excess volatility" puzzle

In this section we discuss another important empirical regularity, which has been called the "excess volatility" puzzle, i.e. the volatility of the exchange rate by far exceeds the volatility of the underlying economic variables. Baxter and Stockman (1989) and Flood and Rose (1995) found that while the movements from fixed to flexible exchange rates led to a dramatic increase in the volatility of the exchange rate no such increase could be detected in the volatility of the underlying economic variables. This contradicted the 'news' models that predicted that the volatility of the exchange rate can only increase when the variability of the underlying fundamental variables increases (see Obstfeld and Rogoff (1996) for a recent formulation of this model).

[6]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 .

In order to deal with this puzzle we compute the noise to signal ratio in the simulated exchange rate. We derive this noise to signal ratio as follows:

$$var(s) = var(f) + var(n) \tag{19}$$

where var(s) is the variance of the simulated exchange rate, var(f) is the variance of the fundamental and var(n) is the residual variance (noise) produced by the non-linear speculative dynamics which is uncorrelated with var(f). Rewriting (19) we obtain

$$\frac{var(n)}{var(f)} = \frac{var(s)}{var(f)} - 1 \tag{20}$$

The ratio $var(n)/var(f)$ can be interpreted as the noise to signal ratio. It gives a measure of how large the noise produced by the non-linear dynamics is with respect to the exogenous volatility of the fundamental exchange rate. We simulate this noise to signal ratio for different values of the extrapolation parameter beta (see figure 14). In addition, since this ratio is sensitive to the time interval over which it is computed we checked how it changes depending on the length of the time interval. In particular, we expect that the noise-to-signal ratio is larger when it is computed on a short than on a long time horizon. We show the results in figure 15.

First, we find that with increasing β the noise to signal ratio increases. This implies that when the chartists increase the degree with which they extrapolate the past exchange rate movements, the noise in the exchange rate, which is unrelated to fundamentals, increases. Thus, the signal about the fundamentals that we can extract from the exchange rate becomes more clouded when the chartists extrapolate more. Second, we find that when the time horizon increases the noise-to-signal ratio declines. This is so because over long time horizons most of the volatility of the exchange rate is due to the fundamentals' volatility and very little to the endogenous noise. In contrast, over short time horizons the endogenous volatility is predominant and the signal that comes from the fundamentals is weak. This is consistent with the empirical finding concerning misalignments we discussed before.

It is also important to relate the noise-to-signal ratio to transaction costs. Therefore we show in figure 16 how the noise-to-signal ratio changes with the size of transaction costs.

We observe that the noise-to-signal ratio increases significantly with the size of transaction costs. An interpretation of this result is that as transaction costs increases the mean-reverting force from fundamentalists is weak while the chartists'force is strong. Thus the noise created by chartists increases and clouds the signal coming from fundamentals.

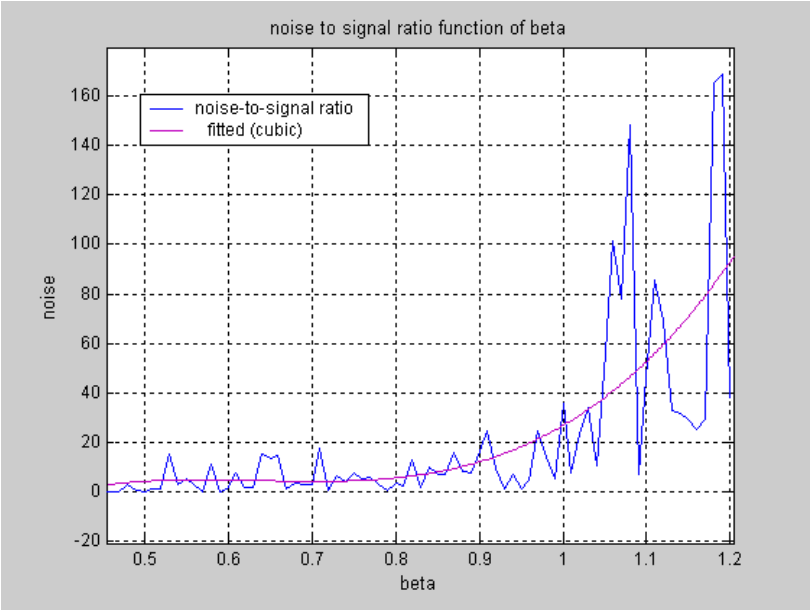


Figure 14:

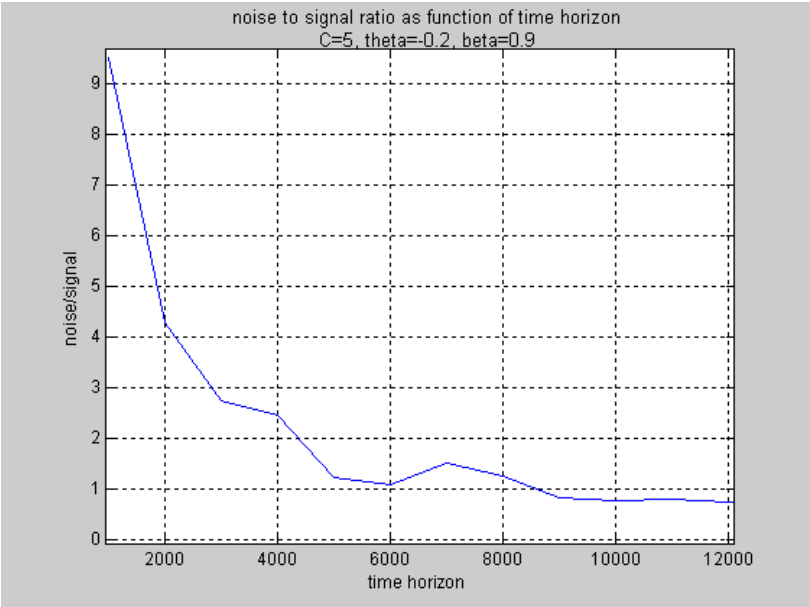


Figure 15:

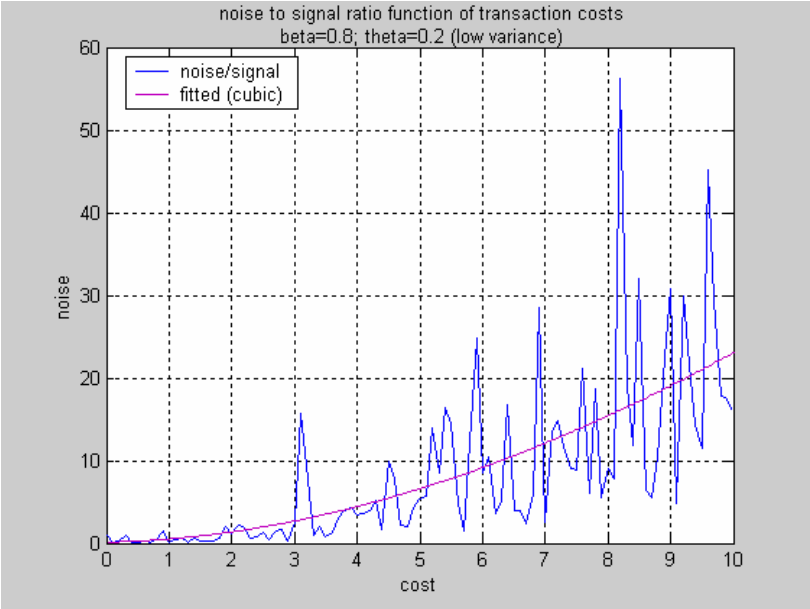


Figure 16:

7.3 Fat tails

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 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, Stork and de Vries (1992), Huisman, et al.(2002)). It implies that most of the time the exchange rate movements are relatively small but that occasionally periods of turbulence occur with relatively large exchange rate changes. However, it has also been detected that the kurtosis is reduced under time aggregation. This phenomenon has been observed for most exchange rates (Lux(1998), Calvet and Fisher(2002)). We checked whether this is also the case with the simulated exchange rate changes in our model.

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 4 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

		Median Hill index (4 samples 2000 observations)		
Parameter values	**Kurtosis**	**2.5% tail**	**5% tail**	**10%tail**
c=5, beta=0.5, theta=0.2	185.4	2.1	2.8	3.1
c=5, beta=0.7, theta=0.2	66.7	1.7	1.6	1.9
c=5, beta=0.7, theta=0.05	5.9	4.6	4.3	4.0
c=5 , beta=0.5, theta=0.1	48.8	2.9	3.4	3.4
c=5 , beta=0.5, theta=0.05	8.5	4.2	4.2	4.0

Figure 17:

of the kurtosis and of the Hill index in figure **??**. 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 figure **??** we show the results for some sets of parameter values[7]. This suggests 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).

7.4 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 Kirman (2002), Lux and Marchesi (2000)). In order to check if our model is capable of reproducing this statistical property we tested if the simulated exchange rate returns have a GARCH structure. We first compute the autocorrelation function of the absolute returns of the simulated exchange rate for different parameter values. In figure 19 we show the autocorrelation function for a particular set of parameters. We find that the autocorrelation parameters

[7]Another empirical regularity of the distribution of exchange returns is its symmetry. We computed the skewness, and we could not reject that the distribution is symmetric.

Kurtosis under time-aggregation

Parameter values	1 period returns	5 periods returns	10 periods returns	25 periods returns
C=5, theta=0.2, beta=1	7.1	4.5	4.0	2.2
C=5, theta=0.2, beta=0.8	40.3	13.3	10.3	4.6
C=5, theta=0.2, beta=0.5	182.9	37.9	24.2	9.9
C=5, theta=0.1, beta=1	3.8	3.2	3.3	2.4
C=5, theta=0.3, beta=1	8.9	4.9	3.6	2.2

Figure 18:

die out slowly. This implies that volatility in the exchange rate returns has a long memory.Moreover, we performed an ARCH test on the residuals of the simulated exchange rate returns and we rejected the null hypothesis of homoskedasticity. Then, we tested for GARCH effects in the exchange rate returns. In order to do so, we chose the simplest possible GARCH specification, i.e. GARCH (1,1)):

$$\Delta s_t = a + \epsilon_t$$

$$\sigma_t^2 = b + \alpha \epsilon_{t-1}^2 + \delta \sigma_{t-1}^2$$

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 figure **??** for different values of the extrapolation parameter β. Although the appropriate orders of lags in the GARCH specification could be identified by the Box-Jenkins methodology applied on the squared residuals ϵ_t^2, we will see that our particularly simple specification fits the empirical evidence rather well.

We observe that the GARCH coefficients, α and δ, are significantly different from zero implying that there is volatility clustering in the exchange rate returns. In addition we find that when for values of β close to one the sum of the α and δ, which is a measure of the degree of the inertia of the volatility, is close to one suggesting 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.

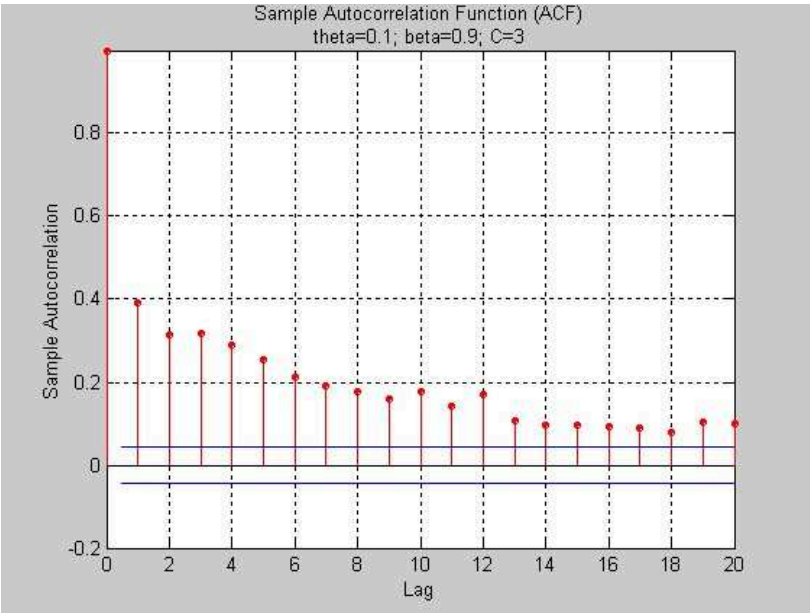


Figure 19:

8 Large and small shocks

In linear models the size of the shocks does not affect the nature of the dynamics. In non-linear models things are different. The size of the shocks matters. This is also the case in our exchange rate model. In order to illustrate this, we simulated the model under two different assumptions about the variance of the shocks in the fundamental exchange rate. In the first case we assume low variance of these shocks, in the second case we assume a high variance (ten times higher). The results of our simulations are presented in figures 21-24. (The simulations shown here are representative for a wide range of parameter values).

Two conclusions follow from a comparison of the low and high variance cases. First, in the low variance case we observe sustained deviations from the equilibrium exchange rate; this is not the case when the equilibrium exchange rate is subject to large shocks (compare figures 21 and 23). Second, the sensitivity to small changes in initial conditions is clearly visible when the variance of the exchange rate is low (see figure 22). When this variance is high, no such sensitivity can be observed (figure 24). It is important to stress that the transactions cost band is the same in both cases. Thus, when the shocks are small relative to the given band of transactions costs, the movements of the exchange rate show more complexity than when the shocks are large.

The previous results are confirmed by a cointegration analysis like the one we performed in section 7.1 where this analysis refers to a low variance environment.

Parameter set: C=5, beta=0.9, theta= 0.1		
	coefficient	*z-statistic*
a	0.019 (0.003)	7.67
b	0.006 (0.0005)	11.62
α	0.65 (0.03)	19.76
δ	0.30 (0.027)	11.05
Log Likelihood Value	*768.1*	
Parameter set: C=5, beta=0.95, theta= 0.1		
	coefficient	*z-statistic*
a	0.019 (0.003)	6.95
b	0.007 (0.0008)	8.68
α	0.57 (0.03)	17.61
δ	0.36 (0.033)	11.35
Log Likelihood Value	*284.1*	
Parameter set: C=3, beta=0.9, theta= 0.2		
	coefficient	*z-statistic*
a	0.004 (0.003)	1.31
b	0.004 (0.0005)	8.0
α	0.30 (0.04)	8.5
δ	0.56 (0.04)	14.85
Log Likelihood Value	*894.5*	

Figure 20: Estimation of GARCH(1,1) on simulated returns

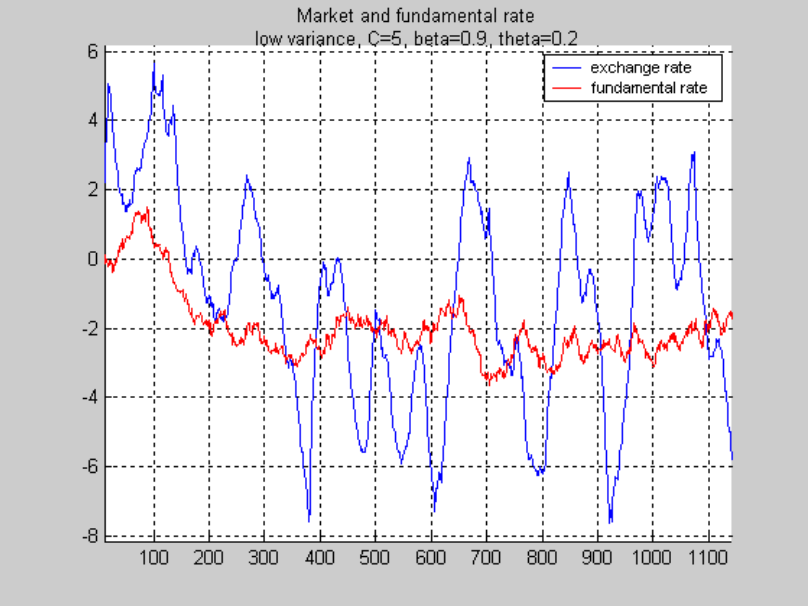


Figure 21:

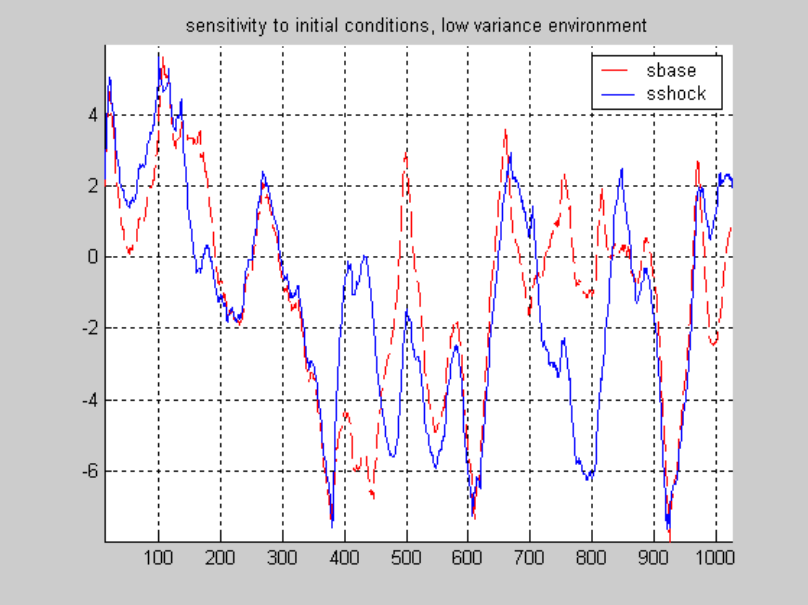


Figure 22:

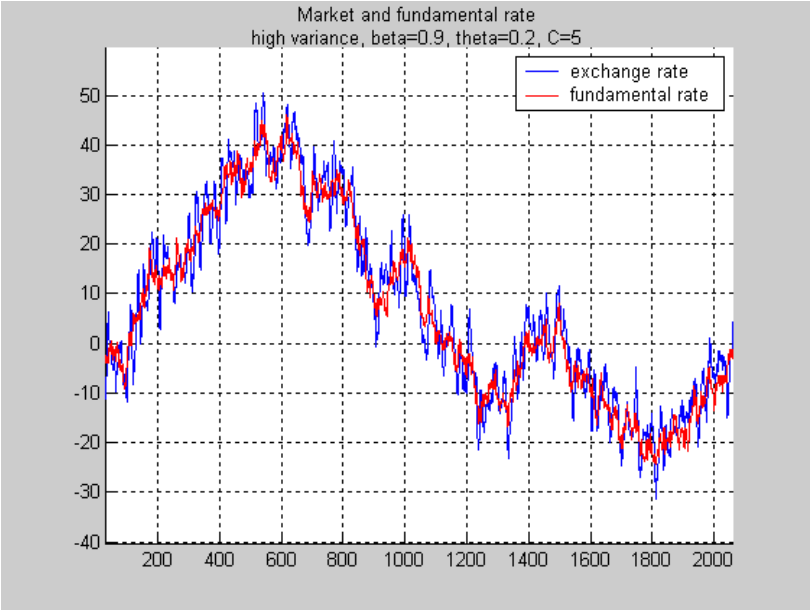


Figure 23:

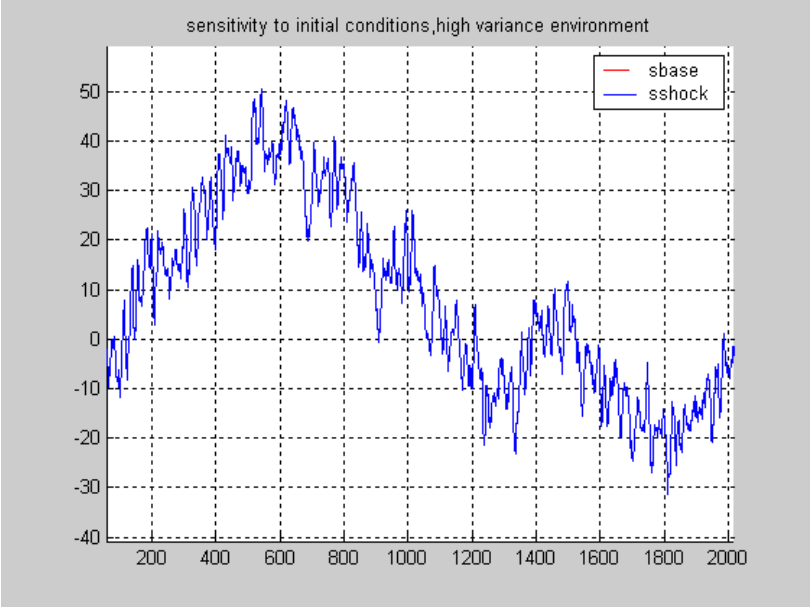


Figure 24:

We show the results for the high variance regime in figure **??** where we contrast the low and the high variance regimes. First the error correction coefficient is much larger in the high variance regime than in the low variance regime. In particular 23% of the adjustment is realised each period in the high variance regime in contrast with only 0.5% in the low variance case. Second, the impact of past changes in fundamentals is significantly higher in the high variance than in the low variance case[8].

The intuition of this result is that when the fundamental shocks are small the exchange rate regularly switches from the dynamics inherent within the transactions cost band to the one prevalent outside the band. This non-linearity produces a lot of noise and complexity in the dynamics of the exchange rate. When the shocks are large relative to transactions cost band the dynamics outside the band mostly prevails, leading to a tighter link between the exchange rate and the fundamental. This feature has also been found to hold empirically. In particular, it has been found that the PPP-relationship holds much tighter in high inflation countries than in low inflation countries (See De Grauwe and Grimaldi(2001)). Put differently, in high inflation countries the link between the exchange rate and one of its most important fundamentals is tighter than in low inflation countries.

It is also interesting to analyse how the variance regimes affect the noise-to-signal ratio. We show the results in figures 27 and 28 where we plot the noise-to-signal ratio against the extrapolation parameter β.

In a low variance environment the noise-to-signal ratio is very large exceeding one by a wide margin. Conversely in a high variance environment the noise-to-signal ratio is below one for all values of β. This confirmes our previuos results, i.e. in high variance environments most of the volatility of the exchange rate is driven by the volatility of fundamentals, while in the low variance environment the exchange rate volatility is due mostly to pure noise.

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 (Taylor and Allen (1992), Wei and Kim (1997)) .

We investigate this issue by analysing how chartism evolves under different conditions. In figure 29 we show the share of chartists in the market for increasing values of the extrapolation parameter β. We obtained the chartists weights by simulating the model over 10000 periods and computing the weight of the last period. Our first finding is that chartism does not disappear, i.e.

[8] As in the low variance case we also performed a cointegration analysis over shorter sample periods. The results contrast with the low variance case. For sample periods of 1000 we find that exchange rate and its fundamentals are cointegrated, while we do not find cointegration in the low variance case.

Figure 25: Parameter estimates of EC model (equation 19) with small and large variance of shocks. (numbers in itlaics are t-ratios; R-square=0.11 (low variance) ans R-square=0.18 (high variance)

Error correction term		Δs_{t-i}				Δs^*_{t-i}			
Small variance regime									
μ	γ	λ_1	λ_2	λ_3	λ_4	φ_1	φ_2	φ_3	φ_4
-0.005	1.06	0.18	0.15	0.09	0.05	0.04	0.03	0.02	0.01
-6. 4	*11.3*	*16.4*	*13.1*	*7.06*	*4.3*	*2.2*	*1.6*	*1.3*	*1.06*
Large variance regime									
μ	γ	λ_1	λ_2	λ_3	λ_4	φ_1	φ_2	φ_3	φ_4
-0.23	0.99	0.08	0.12	0.08	0.05	0.11	0.06	0.02	0.03
-31.3	*356.9*	*7.1*	*11.3*	*7.2*	*5.1*	*5.9*	*3.1*	*0.9*	*1.8*

Figure 26:

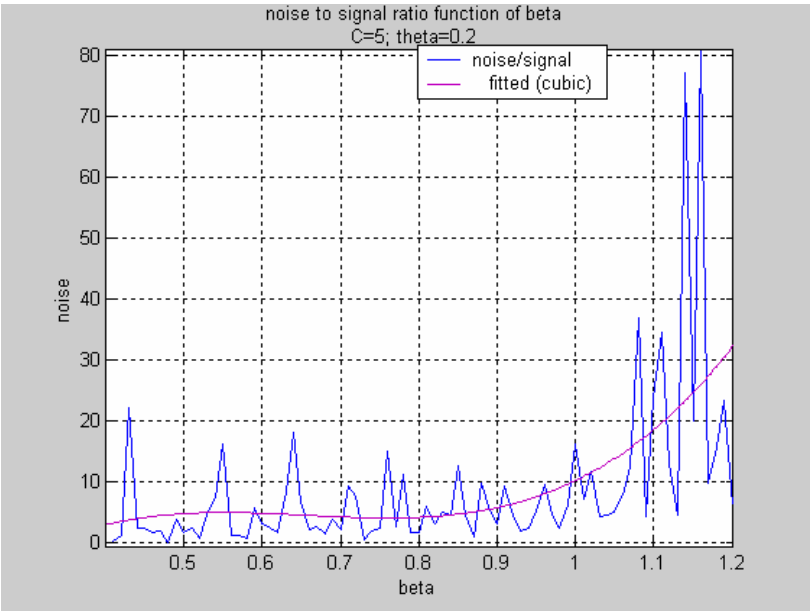


Figure 27:

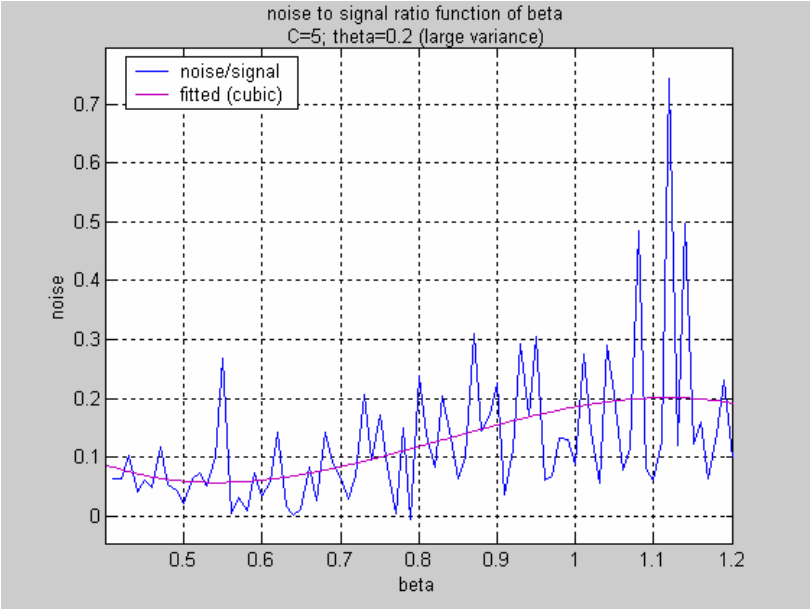


Figure 28:

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 in the market fluctuates around a market share, which exceeds 50%. Thus, in all our simulations of the model we find that chartist rules tend to dominate the fundamentalist rules. (In appendix we show simulations where the weight of chartists is made to vary with the speed of adjustment in the goods market). Third, as the extrapolation parameter β increases the weight of chartists first increases slightly. When β reaches a critical point around 0.9 (which is also the point where the dynamics switches into chaos), the market share of chartists jumps up to approximately 2/3. It then moves down to settle around 64%. This suggest that as chartists become more aggressive in extrapolating past movements of the exchange rate, the chartist rule becomes more profitable attracting more agents to become chartists. Too much extrapolation in turn reduces the relative profitability of chartism and reduces its market share. The latter result suggests that when chartists extrapolate too much the ensuing deviations from the fundamental exchange rate makes fundamentalist forecasting rules more profitable again.

The next step in the analysis consisted in doing stochastic simulations. We simulate two regimes, a low variance and a high variance regime. These two regimes have been defined in the previous section. We show the results in figures 30 and 31. We find that in the low variance regime the share of chartists in the market tends to be higher than in high variance regimes. This result is very

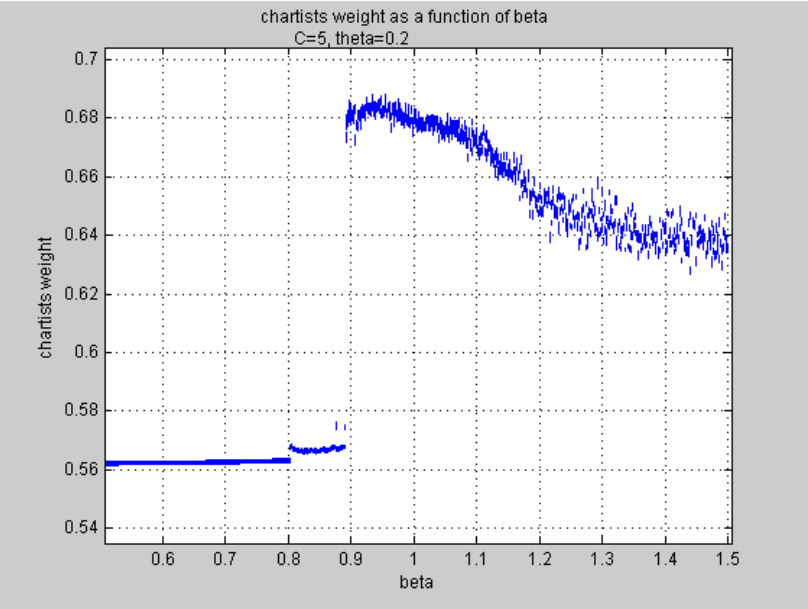


Figure 29:

much in line with the results discussed in the previous section. In a high variance regime the noise to signal ratio of exchange rate movements is relatively low. As a result, movements in the exchange rates convey more information about the underlying fundamental, making fundamentalists rules relatively less risky and thus increasing the weight of fundamentalists (reducing the weight of chartists). Conversely, in a low variance regime, the noise to signal ratio is large, so that exchange rate movements convey less information about the fundamentals. This makes fundamentalist forecasting less precise, and thus riskier. At the same time, the high noise content makes chartists rules relatively profitable attracting more chartists in the market.

A final step in the analysis consisted in analysing the question of how the profits of the chartists and the fundamentalists evolve when the share of the chartists increase exogenously. In order to answer this question we simulated a different version of the the model assuming that the shares of chartists vary exogenously from 0 to 0.9 (instead of being determined endogenously by relative profitability). We then computed the average profits of chartists rules as the share of the chartists increases. We did the same for the fundamentalists profits. We show the results in figures 32 and 33. The results are quite striking. As the weight of chartists increases the profits they make increases strongly. In fact the increase becomes exponential as the share of the chartists in the market becomes very high. In contrast, the profitability of fundamentalists rules is is relatively insensitive to the number of chartists as long as these

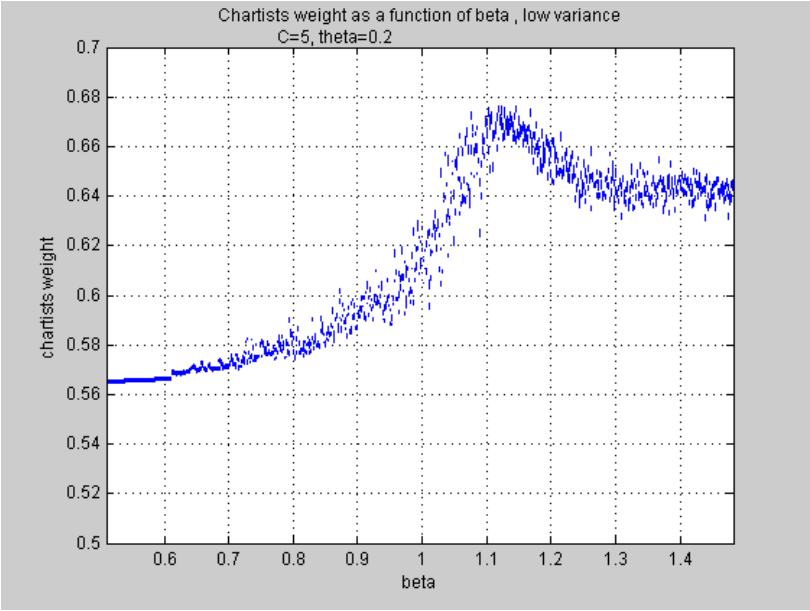


Figure 30:

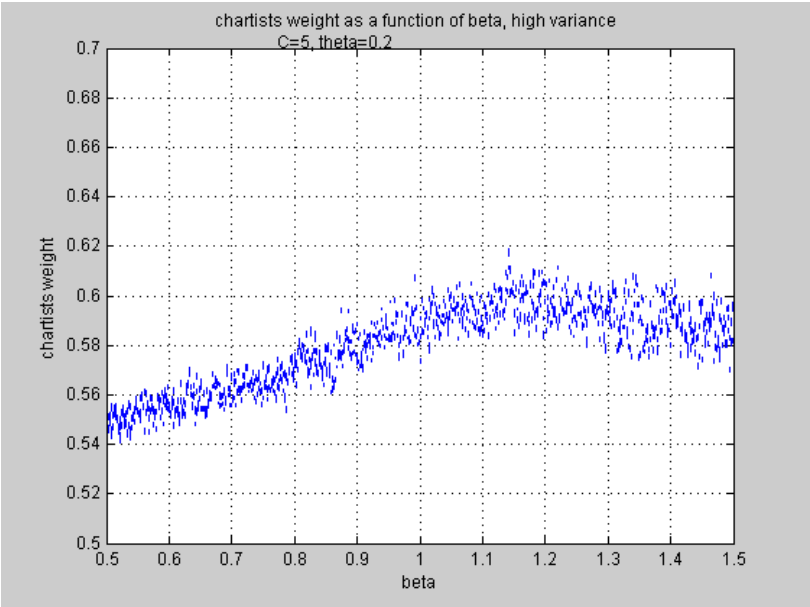


Figure 31:

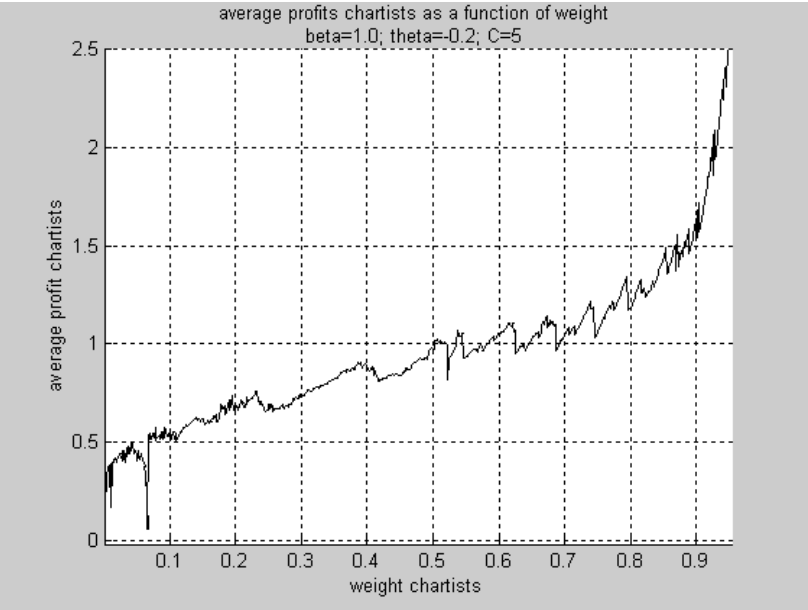


Figure 32:

are not too numerous. All this suggests that there is a selffulfilling evolutionary dynamics present in the system which can be described as follows. As the chartists increase in numbers, their profits tend to increase exponentially. The reason is that the increasing importance of chartists increases the noise in the exchange rate movements, making chartism more profitable. At the same time, the chartists have the effect of "creating smoke around the fundamentals", making fundamentalists' forecasting riskier. Another way to interpret this result is that chartism creates noisy information that becomes the source of profitable speculation. The more chartists there are the more such information is created and the more profitable chartists forecasting becomes. Thus, chartists create an informational environment which makes it rational to use chartists' rules.

Why doesn't all this not lead to a corner solution, i.e. a situation in which chartism drives out all fundamentalists? As we have seen in the previous paragraphs, the share of the chartists in the market is not driven to 1; it always settles below 1. The reason has to do with risk. When the weight of chartists increases in the market, so does volatility. We show this in figure 34 where we present the exchange rate as a function of the weight of chartists. We observe that an increase in the chartists weight leads to a significant increase in volatility especially when the weight approaches 1. Thus, as the weight of chartists in the market increases both profitability and risk of using chartist rules increase. The increasing risk is strong enough to prevents the chartists from completely driving out the fundamentalists and taking over the market.

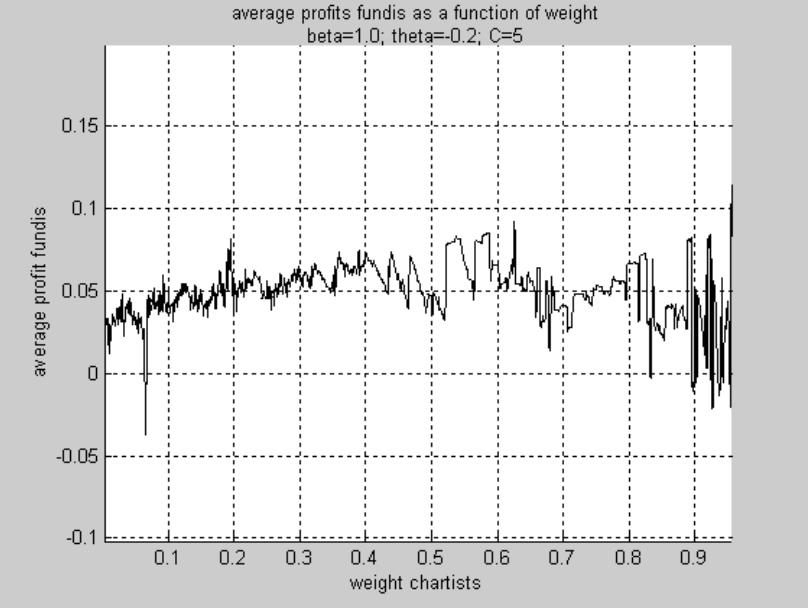


Figure 33:

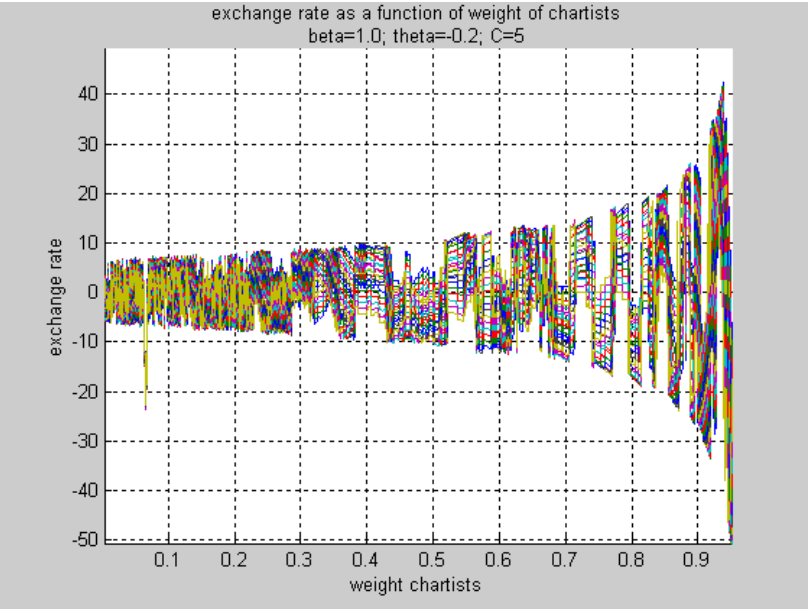


Figure 34:

10 Conclusion

In this paper we use a mean variance optimisation framework to develop a simple non linear exchange rate model with transactions costs and with heterogeneous agents. Transactions costs produce an important non-linearity in the model. Agents are heterogenous in that they have different beliefs, i.e. they use two different forecasting rules. The first type of agetns (the chartists) extrapolate past exchange rate movements. They use a positive feedback rule, which tends to destabilise the market. The second type of agents (the fundamentalists) expect the exchange rate to return to its fundamental value. They use a negative feedback rule, which is stabilising. The relative importance of these two types of agents is assumed to be driven by the relative profitability of their forecasting rules and by the risk associated with the use of such rules. This simple model is capable of creating a remarkably complex dynamics.

The model generates a multitude of fixed-point attractors depending on the initial conditions. Put differently, for each initial condition there is a unique solution. By adding exogenous noise the model produces a complex dynamics that resembles a chaotic dynamics, 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 has a complex structure that might even be chaotic. This implies that the effect of a permanent shock in the fundamentals 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. We also find that the effect of such a shock depends on the exact timing of its occurrence. Thus, history matters. The market has a memory. This constrasts with the prevailing exchange rate models based on the efficient market and rational expectations assumptions that tend to be a-historical.

The empirical relevance of the model is a measure of its quality. Therefore, we analyse to what extent our model is capable of reproducing the exchange rate puzzles that we observe in reality. The first puzzle we analyse 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 of the complex dynamics produced by the interactions between agents using different pieces of information.

Closely related to the disconnect puzzle is the presence of excess volatility of the exchange rate compared to the volatility of its fundamentals. This feature has been widely documented in the empirical literature. Our model mimicks this feature. We find that it is connected to the number of chartists in the market, i.e. the greater the share of chartists the larger is the noise to signal ratio in exchange rate movements.

Third, fat tails and excess kurtosis, which have been detected in the exchange rate returns, are generated by our model. In other words, our model produces a complex dynamics of the exchange rate with intermittency of periods of high

and low turbulence. We find that this alternation of periods of tranquility and turbulence is itself unpredictible.

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

Fifth, the empirical evidence suggests that in environments with high variance of the fundamentals (e.g. in high inflation countries) the link between exchange rate changes and its fundamentals (e.g. inflation rates) is tighter than in low variance environments. We also obtain such a result in our model.This also implies that in high variance environments predicting exchange rate changes using fundamental information should be easier than in low variance environments.

Finally, we investigated under what conditions chartism is evolutionary stable.We found that chartism does not disappear, i.e. in all simulations for many different parameters configurations we find that the number of chartists never goes to zero. This result is consistent with the empirical evidence of the importance of chartism in foreign exchange market. We also detected a self-fullfilling character of chartist profitability, i.e.when more chartists enter the market they create more noise and thereby make chartists rules more profitable, inducing more entry. Another way to interpret this result is that chartism creates noisy information that becomes the source of profitable speculation. The more chartists there are the more such information is created and the more profitable chartists forecasting becomes. Thus, chartists create an informational environment which makes it rational to use chartists' rules.This process is stopped, however, because of increasing risk generated by the increased noise that the same chartists produce.

11 References

Baxter, M., Stockman, 1989, A., "Business Cycles and the Exchange Rate Regime. Some International Evidence", Journal of Monetary Economics, 23, may 377-400.

Brock, W., and 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. and Lai K., 2000. " On the purchasing power parity puzzle". Journal of International Economics, 52 .

Cheung Y., Lai K. and 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.

Copeland, L., 2000, Exchange Rates and International Finance, 3rd ed., Prentice Hall.

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

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

De Grauwe, P., and Grimaldi, M., 2001, "Exchange Rates, Prices and Money: A Long Run Perspective", International Journal of Finance and Economics, 6, no. 4, pp. 289-314.

De Grauwe, P., and Vansteenkiste, I., 2001," Exchange Rates and Fundamentals. A Non-linear Relationship?", CESifo Working Paper, no. 577, October.

de Vries, C., 2000, "Fat tails and the history of the guilder", Tinbergen Magazine, 4, Fall, pp. 3-6.

De Long, J., Bradford, B., Schleiffer and Summers, L., 1990, "Noise Trader Risk in Financial Markets", Journal of Political Economy.

Dornbusch R., 1976, "Expectations and exchange rate dynamics", Journal of Political Economy 84.

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. and Morley J., 2001, "The adjustment of prices and the adjustment of the exchange rate", Discussion paper, Department of Economics, University of Wisconsin.

Evans, M., and Lyons, R., 1999, "Order Flow and Exchange Rate Dynamics", NBER Working Paper, no. 7317.

Fischer S., 1982," Relative prices and inflation in the United States and Germany", European Economic Review, No. 18.

Flood, R, and Rose, A., 1995, "Fixing the Exchange Rate Regime: A virtual Quest for Fundamentals", Journal of Monetary Economics, 36, August, 3-37.

Frankel, J., and Froot, K., 1986, "The Dollar as a Speculative Bubble: A Tale of Fundamentalists and Chartists", NBER Working Paper, no. 1963.

Frenkel J., 1981, "The collapse of purchasing power parity during the 1970s", European Economic Review 7.

Goodhart, C., 1989, "News and the Foreign Exchange Market", LSE Financial Markets Group Discussion paper, 71.

Goodhart, C., and Figliuoli, 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" PhD Thesis K.U.Leuven

Hallwood, P., MacDonald, R., 1994, International Money and Finance, 2nd ed., Blackwell, Oxford.

Huisman, R., Koedijk, K., Kool, C., and Palm, F., 2002, The tail-fatness of FX returns reconsidered, in DE Economist, 150, no. 3, September, 299-312.

Isard, P., 1995, Exchange Rate Economics, Cambridge University Press.

Kilian L. and M. Taylor, 2001, "Why is it So Difficult to Beat the Random Walk Forecast of Exchange Rates?" Mimeo, University of Warwick, pp. 29.

Kurz, M., 1994, "On the Structure and Diversity of Rational Beliefs", Economic Theory, 4, 877-900.

Kurz, M., and Motolese, M., 2000, "Endogenous Uncertainty and Market Volatility", mimeo, Stanford University.

LI K., 1999, "Testing symmetry and proportionality in PPP : A panel data approach", Journal of Business and Economic Statistics 17 (4) , 409-418.

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, vol.33.

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

Lyons, R., 2001, The Microstructure Approach to Exchange Rates, MIT Press, Cambridge, Mass.

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

Meese, R., and 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., and 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. and Rogoff, K., 1996, Foundations of International Macroeconomics, MIT Press, Cambridge, Mass.

Obstfeld, M., and Rogoff, K., 2000, "The Six Major Puzzles in International Macroeconomics: Is there a Common Cause?", NBER Working Paper no. 7777, July.

Okun A., 1971, "The mirage of steady inflation", Brooking papers on Economic activity.

Rogoff, K., 1996, "The purchasing power parity puzzle", Journal of Economic Literature, 34, June, 647-668.

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), pp. 101-121.

Schleiffer, A., 2000, Introduction to Behavioural Finance, Clarendon Press.

Taylor, M., and Allen, H., 1992, "The Use of Technical Analysis in the Foreign Exchange Market", Journal of International Money and Finance, 11, 304-14.

Taylor M., Peel D., and 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 Shang-Jin and Kim Jungshik 1997. "The big players in the foreign exchange market:do they trade on information or noise?". NBER working paper 6256.

Williamson, J., 1985, "The Exchange Rate System", Policy Analyses in International Economics, 5, Institute for International Economics, Washington, D.C.

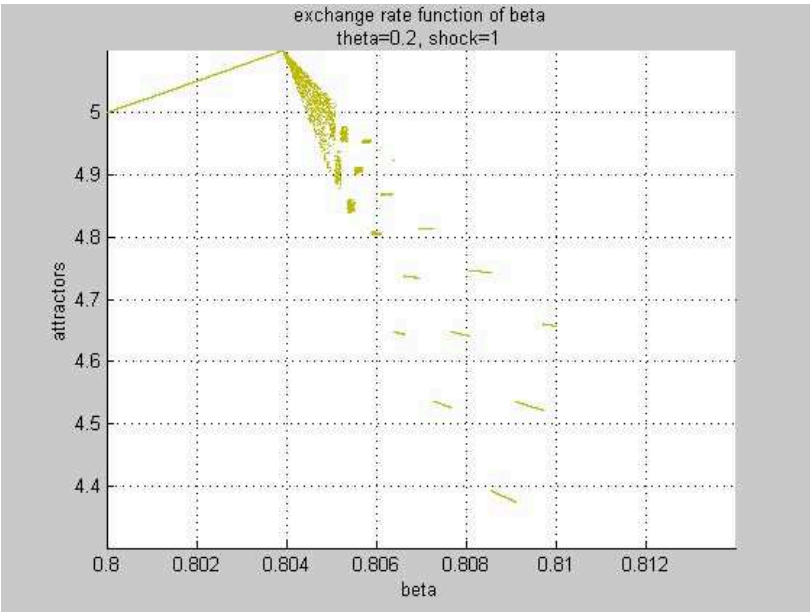


Figure 35:

12 APPENDIX 1: The fractal nature of the fixed point solutions

In this appendix we explore the nature of the critical points where the solution moves from a continuous line into a fractal. Figures 35 and 36 show consecutive enlargements of figure 2. We obtain very similar discontinuities when we enlarge figure 4 which shows the fixed points as a function of different initial conditions. This enlargement is shown inf figure 37. Note the similarity between these diagrams: small changes in a parameter (beta) have the same effect on the solutions as small changes in initial conditions.

13 APPENDIX 2: Additional sensitivity analysis

In this appendix we show the result of a sensitivity analysis where we allow the speed of adjustment in the goods market to change. We construct similar 'bifurcation' diagrams as in the main text, which show how the equilibrium exchange rate is affected by increasing speeds of adjustment in the goods market. In figures 38 to 40 we show three such graphs for three different values of the chartists' extrapolation parameter β. For low values of β, i.e. less than 0.9 (here we show an example for $\beta = 0.8$) we find a fixed point solution. As

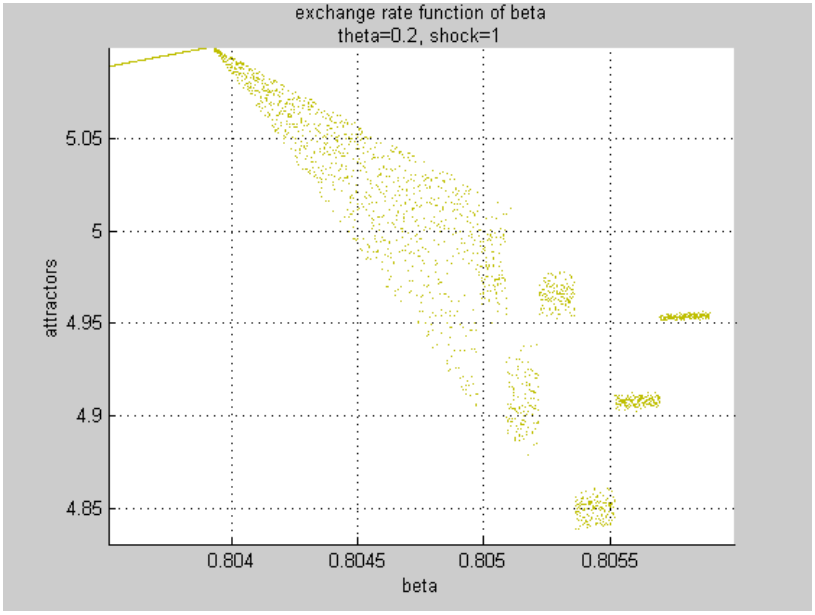


Figure 36:

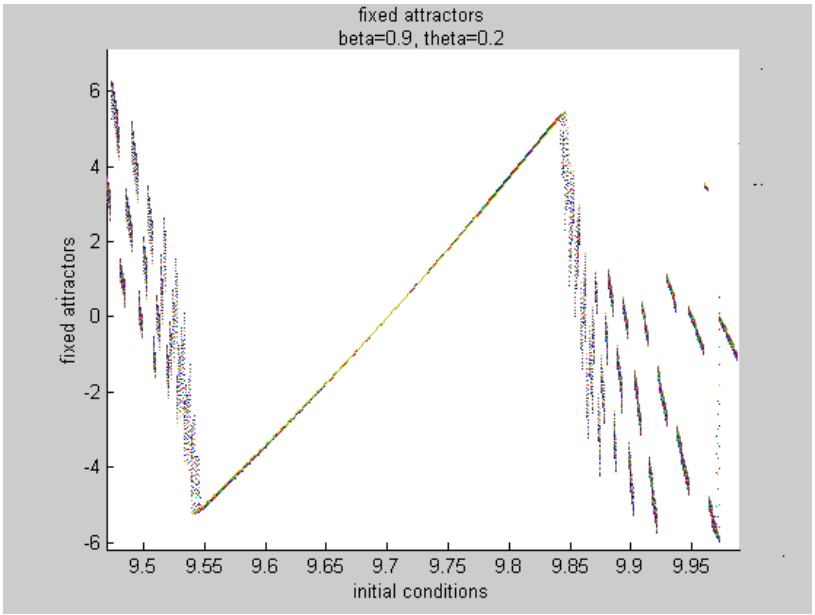


Figure 37:

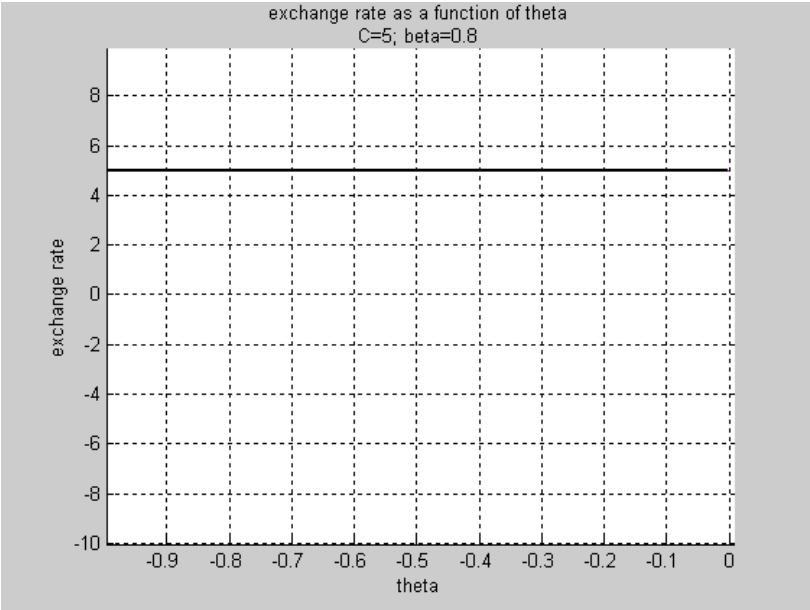


Figure 38:

was shown in the main text the value of this fixed point depends on the initial conditions. It does not depend on the speed of adjustment in the goods market, θ. For $\beta = 0.9$,we find that the solutions switch in an out of the chaotic region depending on the value of θ. Finally for values of β exceeding 0.9, we obtain chaos for all values of θ.

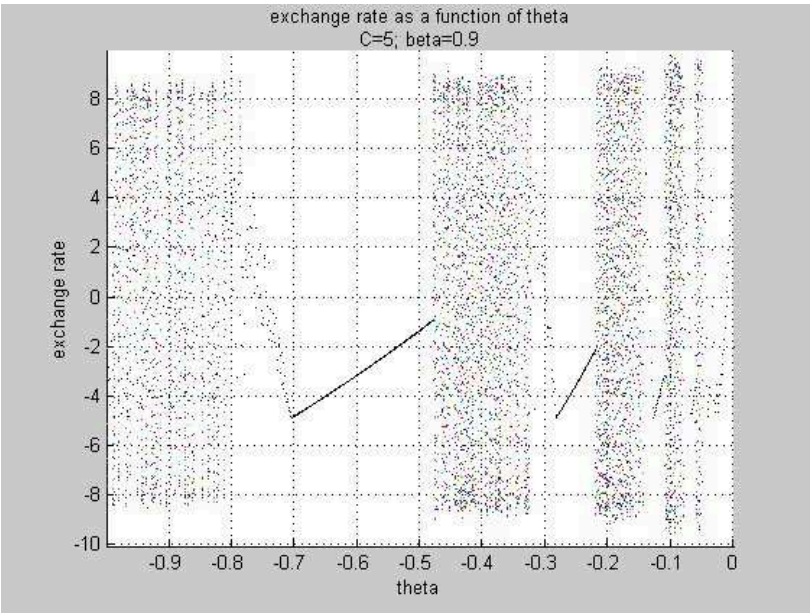


Figure 39:

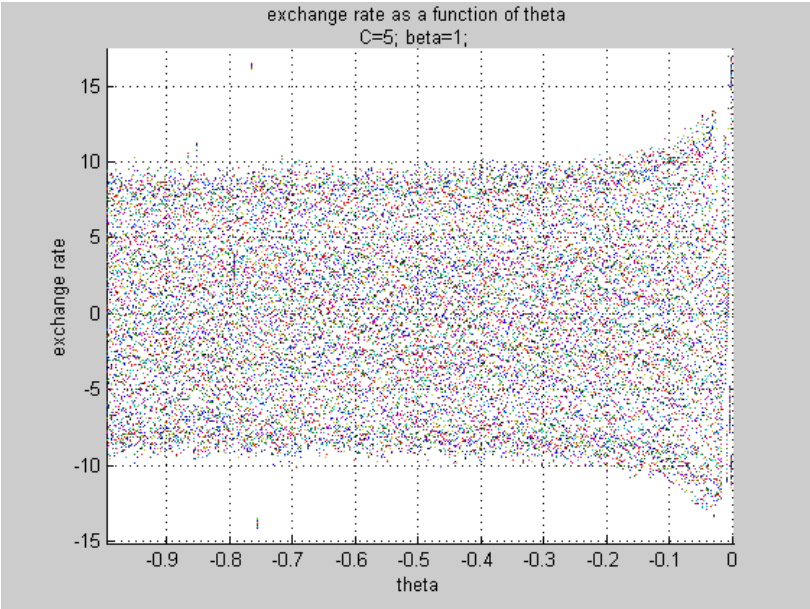


Figure 40: