Exchange rate puzzles.
A tale of switching attractors.

Paul De Grauwe and Marianna Grimaldi[*][†]
University of Leuven

June 12, 2003

Abstract

The rational expectations efficient market model of the exchange rate has failed empirically. In this paper we develop a model of the exchange rate in which agents use simple forecasting rules. Based on an ex post evaluation of the relative profitability of these rules they decide whether to switch or not. In addition, transactions costs in the goods market are introduced. 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 tranquillity and turbulence alternate in unpredictable manner. Finally we show that this model mimics most of the empirical puzzles uncovered in the literature.

JEL classification: F31, F41

Keywords: exchange rate, heterogeneous agents, chartists and fundamentalists, transaction costs

[*]We are grateful for many useful comments to Volker Bohm, Yin-Wong Cheung, Hans Dewachter, Roberto Dieci, Marc Flandreau, Cars Hommes, Thomas Lux, Richard Lyons, Michael Moore, Ronald McDonald, Assaf Razin, Peter Westaway, Philip Lane, Piet Sercu and Peter Sinclair.

[†]Corresponding address: Paul De Grauwe, Naamsestraat 69, 3000 Leuven, Belgium.Phone: +32 16 326836, fax:+3216326796, email:paul.degrauwe@econ.kuleuven.ac.be

1 Introduction

The rational expectations efficient market model developed during the 1970s has dominated our thinking about exchange rates. This model led to the propositions, first, that exchange rate changes can only occur because of unexpected movements (news) in the underlying fundamental economic variables (inflation, growth of output, interest rates, etc.), and, second, that the link between exchange rates and fundamentals is a stable one. Well-known examples of the rational expectation efficient market model is the monetary model, the Dornbusch model (Dornbusch(1976)) and the portfolio balance model. Although these models continue to be popular and maintain a prominent place in textbooks, they have failed empirically. The most notorious empirical rejection was made by Meese and Rogoff at the beginning of the 1980s (Meese and Rogoff(1983)). This led to a large empirical literature that uncovered a number of empirical puzzles concerning the behaviour of the exchange rate, which could not be explained by the 'news' models.

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.

Other empirical anomalies have been uncovered over the years. One is the puzzle of "excess volatility" of the exchange rate, 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)).

Another puzzle is that the distribution of the exchange rate returns is not normal. Most of the empirical findings document that the exchange rate returns have fat tails (see de Vries(2001), Lux T. (1997, 1998), Lux and Marchesi (1999, 2000). This evidence is difficult to rationalise in existing exchange rate models, since there is little evidence of fat tails in the fundamental variables that drive the exchange rate in these models.

The empirical failure of the exchange rate models of the 1970s has led to new attempts to model the exchange rate. These attempts have led to three different modelling approaches. The first one uses the Obstfeld–Rogoff framework of dynamic utility optimisation of a representative agent. This approach although promising is still waiting for empirical confirmation (Obstfeld and Rogoff (1996)).

A second approach starts from the analysis of the microstructure of the foreign exchange market (Evans and Lyons (1999), Lyons (2001)) . This approach has led to new insights into the way information is aggregated and is important for the understanding of the short-term behaviour of the exchange rate.

Finally, a third approach recognises that heterogeneous agents have different beliefs about the behaviour of the exchange rate. These different beliefs introduce non-linear features in the dynamics of the exchange rate (Kirman (1993), Brock and Hommes(1998), Lux (1998)) .

Recently heterogeneity of agents has also been 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.

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 , and by switching to the more profitable rules. In addition, we make use of the recent empirical evidence, which has stressed the importance of transactions costs in the goods market for our understanding of the dynamics of exchange rate adjustments (Michael, Nobay, Peel (1997), O' Connell (1998), Obstfeld and Rogoff (2000), Engel(2000) and Kilian and Taylor (2001)). 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.

The paper is organised as follows. In section 2 we present the theoretical model. In sections 3 ,4, and 5 we analyse its features, while in sections 6 , 7 and 8 we analyse its empirical predictions.We conclude in section 9.

2 The model

In this section we develop a simple non-linear model of the exchange rate . The model consists of three building blocks. First, agents decide the optimal portfolio using a mean-variance utility framework. Second they make forecasts about the future exchange rate based on simple rules. Third, they evaluate these rules ex-post by comparing their risk-adjusted profitability.

2.1 The optimal portfolio

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:

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

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[1]. 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:

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

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 :

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

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

[1] 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.

their respective variances. Thus, when for example the chartists' forecasts have a high variance the weight of the chartists in the determination of the market exchange rate is reduced.

2.2 The forecasting rules

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 and $0 < \psi < \infty$.

In addition, the fundamentalists take the existence of transaction costs in the goods market into account. There is an increasing body of theoretical literature stressing the importance of transactions costs in the goods markets 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 in the goods markets has also been confirmed empirically (Taylor, Peel, and Sarno(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 behave differently depending on whether the exchange rate is within or outside the transaction costs band. 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 forecasting rule as in equation 8. More formally,

when $\mid s_t - s_t^* \mid > C$ holds, then equation 8 applies[2].

However when the exchange rate deviations from the fundamental value are smaller than the transaction costs in the goods markets, there is no mechanism

[2] Note that since $\psi < \infty$ market inefficiencies other than transaction costs continue to play a role when the exchange rate moves outside the transaction costs band. As a result, these inefficiencies prevent the exchange rate from adjusting instantaneously.

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 and the best they can do is to forecast no change. More formally,

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

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,

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

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:

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

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, as the misalignment increases, the fundamentalists become increasingly more confident that the exchange rate will revert to its fundamental value. As a result, their risk perception declines. In contrast the chartists do not take into account the misalignment.

2.3 Fitness of the rules

The next step in our analysis is to specify how agents evaluate the fitness of these two forecasting rules. The general idea that we will follow is that agents use one of the two rules, compare their (risk adjusted) 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 will continue to be followed if they pass some "fitness" test (profitability test).

We start by specifying 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:

$$n_{c,t} = n_{c,t-1} + n_{f,t-1} p_t^{fc} - n_{c,t-1} p_t^{cf} \tag{12}$$

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

where $n_{c,t}$ and $n_{f,t}$ are the number of chartists and fundamentalists in period t; p_t^{cf} represents the fraction of the chartists who decide to become fundamentalists in period t, and p_t^{fc} is the fraction of the fundamentalists who decide to become chartists in period t.

These fractions are assumed to be a function of the profitability of the forecasting rules and the risk associated with their use. The fractions are specified as follows[3]:

$$p_t^{fc} = \frac{exp\left[\gamma \left(\pi_{c,t-1} - \mu \sigma_{c,t-1}^2 \right) \right]}{exp\left[\gamma \left(\pi_{c,t-1} - \mu \sigma_{c,t-1}^2 \right) \right] + exp\left[\gamma \left(\pi_{f,t-1} - \mu \sigma_{f,t-1}^2 \right) \right]} \tag{14}$$

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

$$p_t^{cf} = \frac{exp\left[\gamma\left(\pi_{f,t-1} - \mu\sigma^2_{f,t-1}\right)\right]}{exp\left[\gamma\left(\pi_{c,t-1} - \mu\sigma^2_{c,t-1}\right)\right] + exp\left[\gamma\left(\pi_{f,t-1} - \mu\sigma^2_{f,t-1}\right)\right]} \tag{15}$$

where $\pi_{c,t-1}$and $\pi_{f,t-1}$ are the net profits made by chartists and fundamentalists forecasting the exchange rate in period t-1. They make a profit when they correctly forecast the direction of the exchange rate movement. 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. In the case of the fundamentalists we assume that there is a fixed cost of collecting fundamental information.

Equations 14and 15 can be interpreted as follows. When the risk adjusted profits of the chartists rule increases relative to the risk adjusted profits of the fundamentalists rule, then the fraction of the fundamentalists who become chartists in period t 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 .

3 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. 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 also 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 figures 1 and 2 where we plot the fixed point solutions (attractors) as a function of the different initial conditions. Figure 1 shows the fixed attractors for a low value of γ and figure 2 for a high value of γ. 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 the complex pattern of these fixed point solutions, with many discontinuities. This has the implication that a

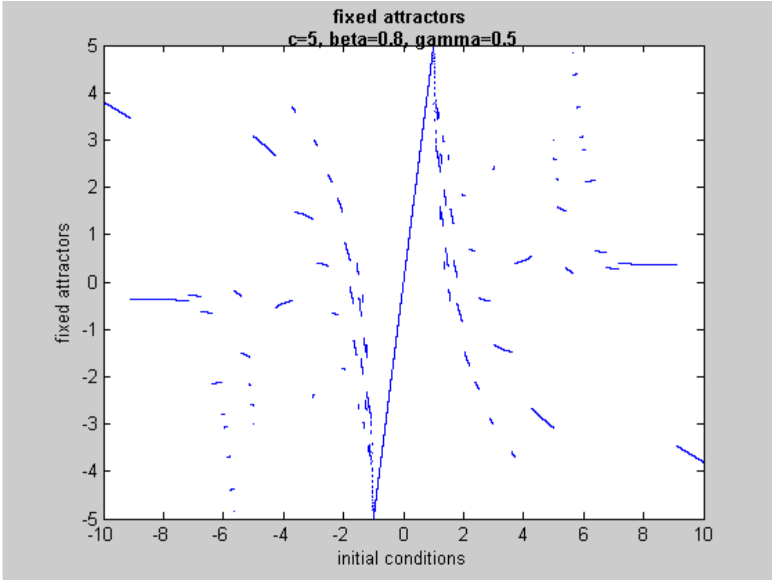

Figure 1:

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 7.

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.

4 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, β,on the

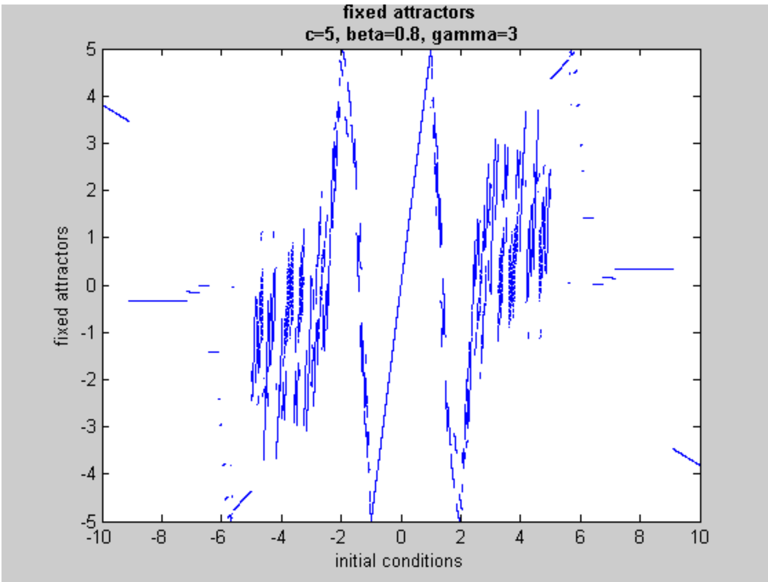

Figure 2:

sensitivity of the switching rule, γ,and on transactions costs.

4.1 Sensitivity with respect to β

Figures 3 and 4 show examples of such diagrams for β. 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 $\beta \leq 0.9$ we obtain a unique fixed point solution for each value of β. 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.

When comparing Figures 3 and 4 we also 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 can be shown that when the system is close to the border between fixed point solutions and chaotic solutions, initial conditions matter a great deal in that they determine whether the system will move to a fixed point or to chaos. We show this feature in Figure 5 where we fix $\beta = 0.9$ and vary the initial conditions. We now observe that the different initial conditions lead to switches in and out of the chaotic region. This feature suggests that there are attractors (some fixed points others strange attractors) located in different basins of attraction. The border line between these different basins is itself complex. As a result, small differences in the intial conditions can lead the system towards different attractors.

These features illustrate the great complexity in the exchange rate dynamics. As 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 complicates the making of standard statistical inferences from the distribution of the exchange rate changes.

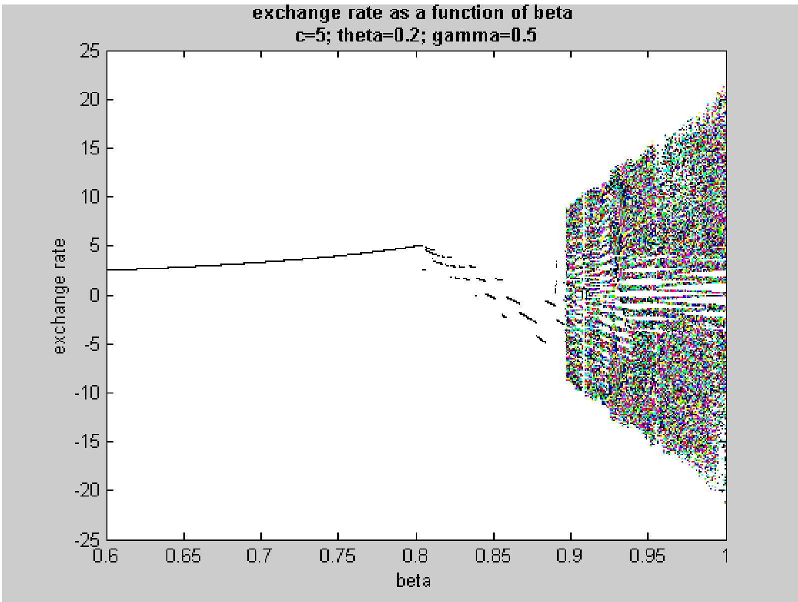

Figure 3:

4.2 Sensitivity with respect to γ

In this section we analyse the sensitivity of the solutions with respect to change in the parameter γ which measures the sensitivity of the switching rules with respect to profits. We show the results in figures 6 and 7 for two different values of the extrapolation parameter, β.When β is sufficiently low ($\beta = 0.8$) we obtain fixed point solutions for all values of γ (figure 6). When $\beta = 0.9$ which as we have seen in the previous section, constitutes the boundary value between fixed point and chaotic solution, variations in γ lead the solutions to switch in and out of chaos (figure 7). Thus changing γ produces similar effects as changing

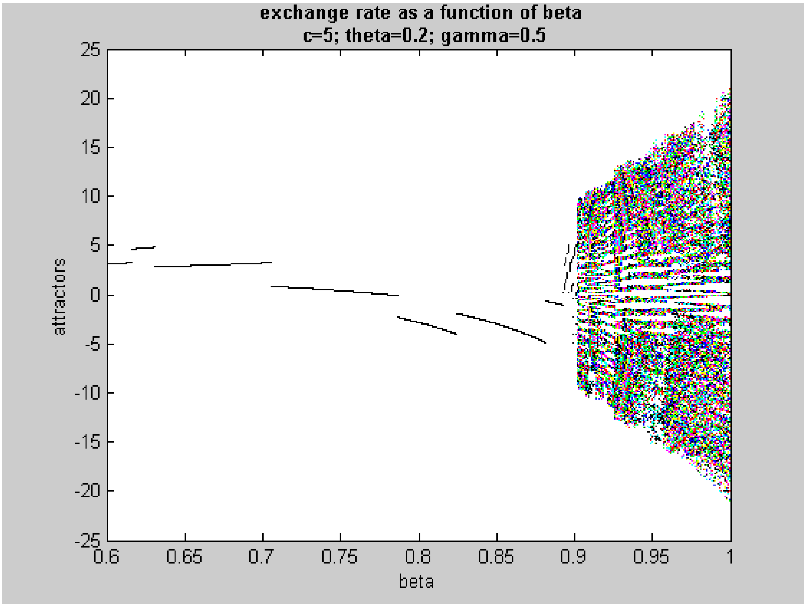

Figure 4:

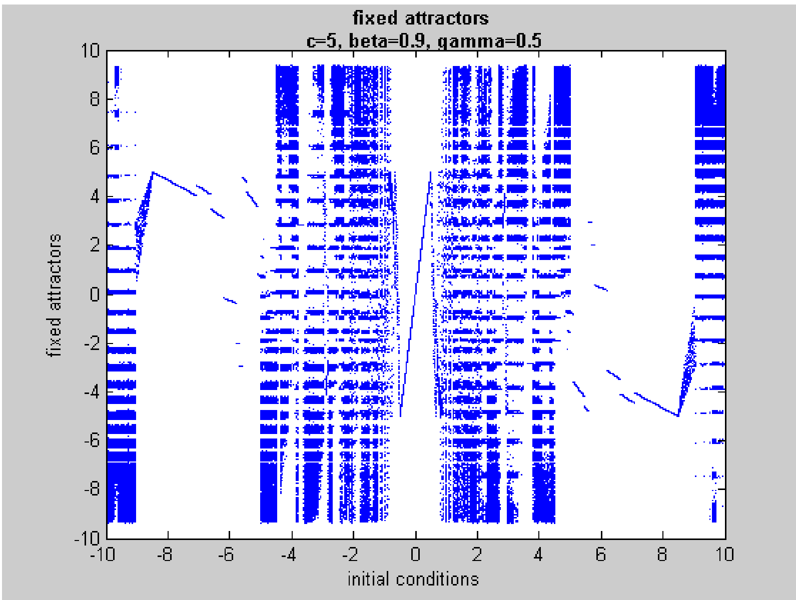

Figure 5:

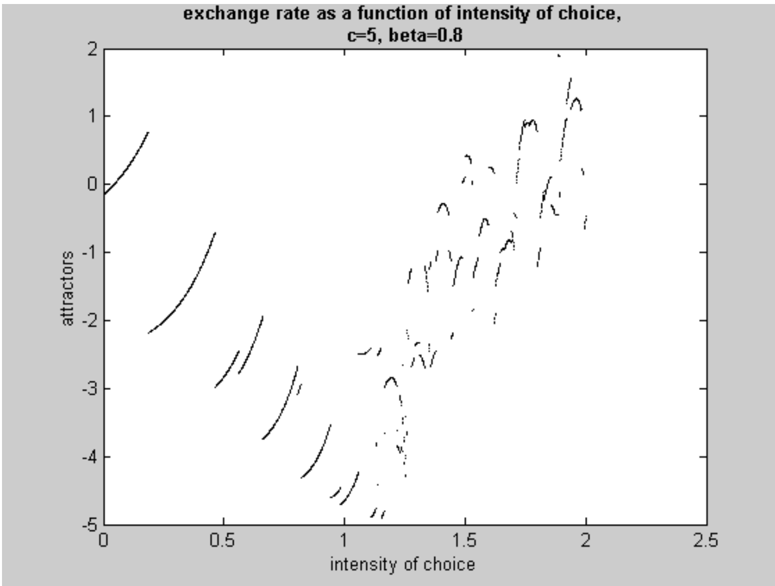

Figure 6:

the initiak conditions. Note, however, that for small values of γ we obtain fixed point solutions.

4.3 Sensitivity with respect to transactions costs

We also investigated the importance of transaction costs. In order to do so, we produced similar bifurcation diagrams as in the previous sections. We now set out the transactions costs on the horizontal axis while we fix β and γ. In figure 8 we set $\beta = 0.8$ and $\gamma = 0.5$. We observe that with a low value of β we obtain fixed point solutions for all values of the transactions costs. Note , however, that when transactions costs are not too high these fixed points are characterised by the same discrete displacements we ahve observed in the previous sections. When transactions costs are sufficiently high, we obtain a unique fixed point for all subsequent values of transactions costs. This has to do with the fact that when transactions costs are large enough the exchange rate stays within the transactions costs band forever. As a result, fundamendalist forecasts disappear from the market and the model becoems linear, producing a unique fixed point.

In figure 9 we show the case of a large extrapolation parameter β. We observe switching in and out of chaos. Note that as transactions costs increase the spread of the possible solutions increases.

It is also important to analyse the dynamics of the weights of chartists and

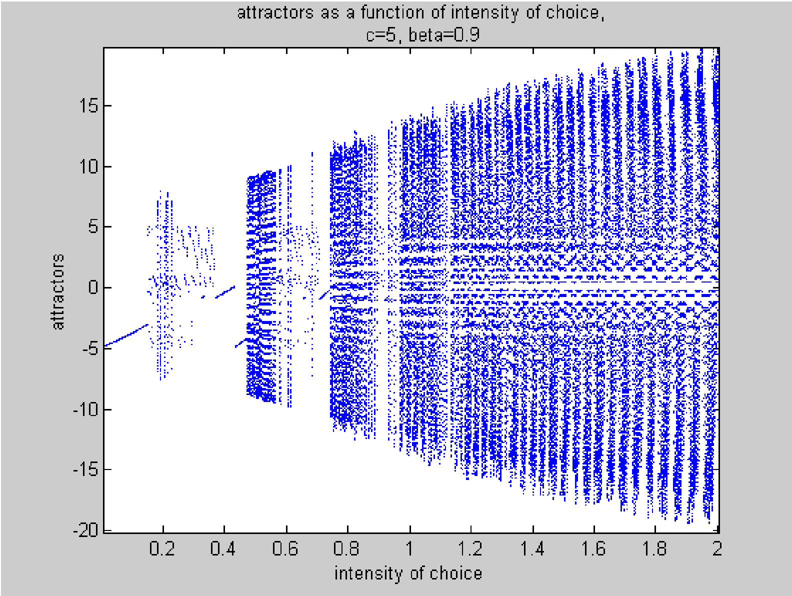

Figure 7:

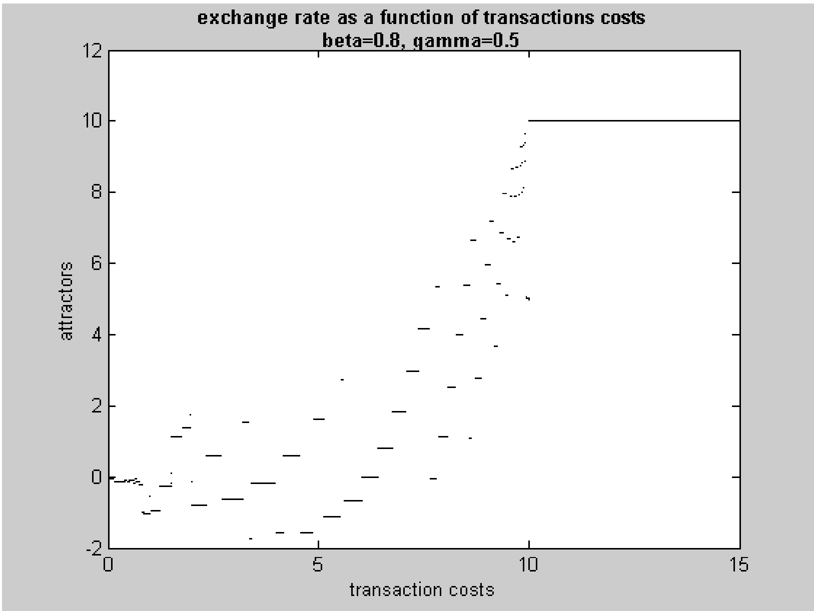

Figure 8:

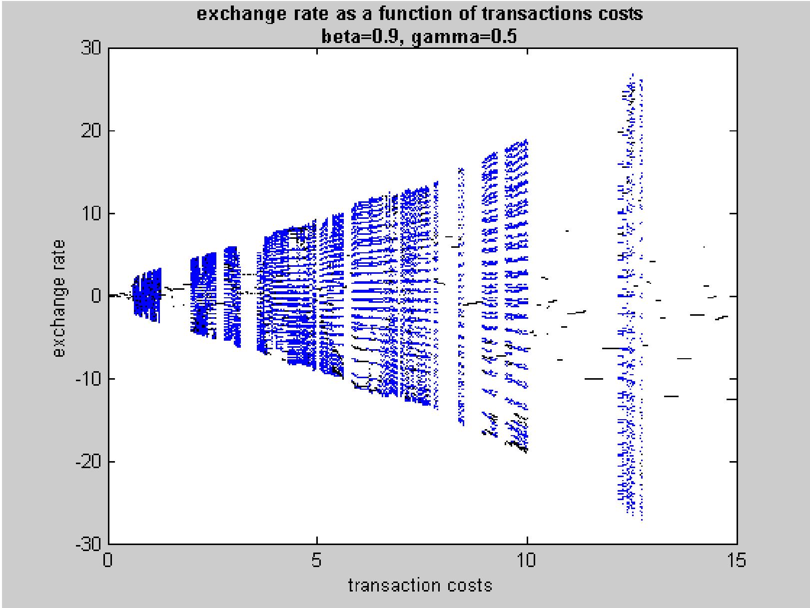

Figure 9:

fundamentalists. We find that each time 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 rate reaches a fixed-point solution, chartists and fundamentalists expect no change anymore. 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 15 and 14) assures that they will be equally represented in the market. When, however, the exchange rate is in a chaotic region, the weights of chartists move within the full interval between 0 and 1. However, in the chaotic region, the chartists clearly dominate. We present an example of this phenomenon in figure 10 where we show the weight of the chartists as a function of γ. (Note that figure 10 should be read in conjunction with figure 7 which presents the exchange rate attractors as a function of γ). In section 8 we will return to the issue of the profitability of chartists' rules 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)). 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. In addition

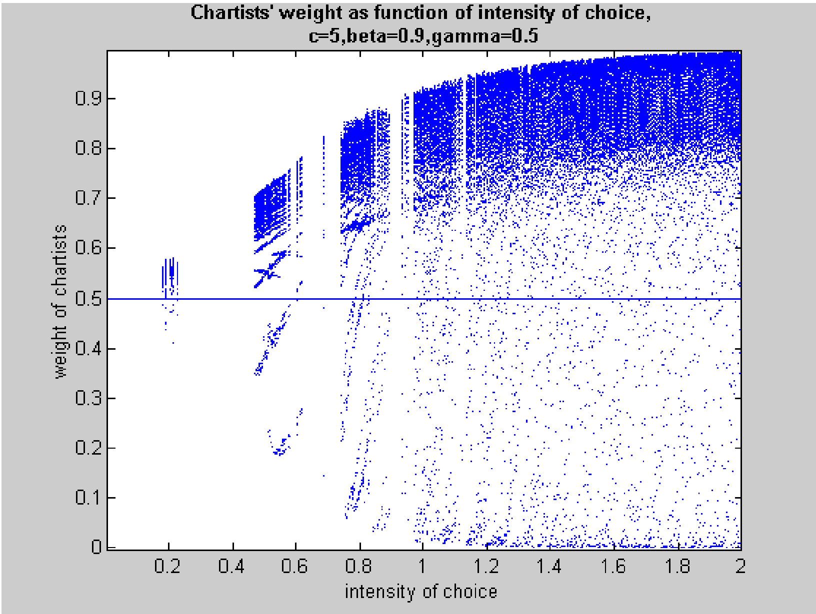

Figure 10:

in section 6 we will calibrate the model in such a way that it reproduces the main statistical properties of exchange rate movements. It will be shown that the parameters that mimick these statisitical properties best do not produce deterministic chaos.

5 The stochastic version of the model

We now introduce stochastic disturbances to the model. In our model these disturbances appear in that we assume that the fundamental behaves as a random walk. We simulate the model with a combination of parameter values that we refer to as the "standard case". This includes setting c=5, β= 0.9 and θ= 0.2. and $\gamma = 0.5$. (Similar results are obtained for a wide range of parameter values. In addition it will be shown in section 6 that these parameter values reproduce the statistical properties observed in exchange rate movements).

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 driving the fundamental. We show the time paths of the (market) exchange rate in figure 11.

We observe that after a certain number of periods the two exchange rates

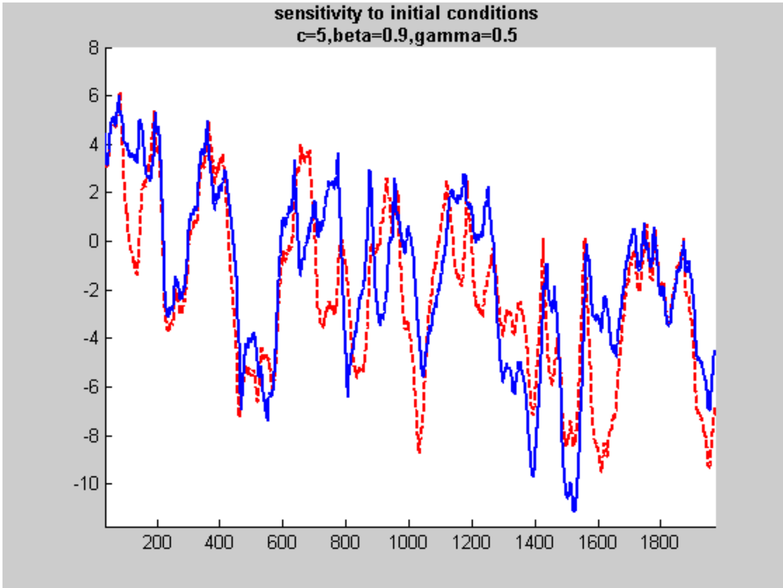

Figure 11:

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 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 they are attracted by fixed- points that are located in different basins of attraction. 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 located in different basins of attraction 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 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 12 where we show the effect of the same permanent shock of 10 in the fundamental rate on the exchange rate.

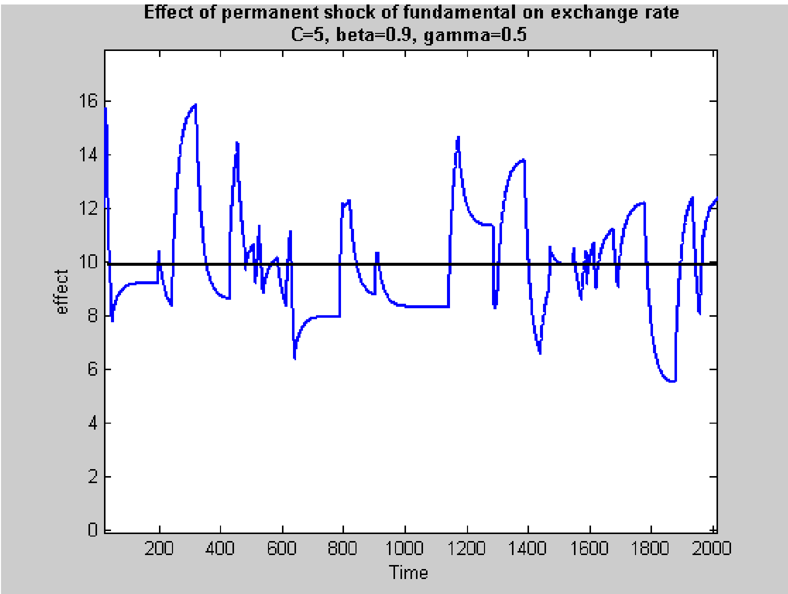

Figure 12:

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.

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 13.

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

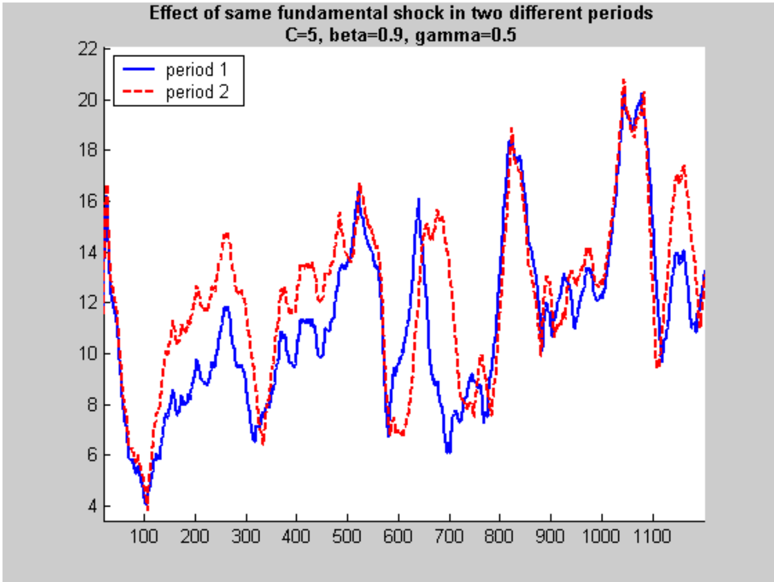

Figure 13:

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 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. This difference is due, among others, to the fact that at each point in time there is a different composition of chartists and fundamentalists in the market, which itself is due to different past performances of chartists and fundamentalists forecasting rules. As a result, the same fundamental shock applied at different time periods is "perceived" differently in the market, e.g. at one moment there are fewer fundamentalists than at another moment so that the same fundamental shock gets less attention. Thus, initial conditions (history) matters to forecast the effect of shocks. Since each initial condition is unique, it becomes impossible to forecast the precise effect of a shock at any given point in time.

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.

6 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'.

6.1 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[4]. 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 14 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

[4] 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.

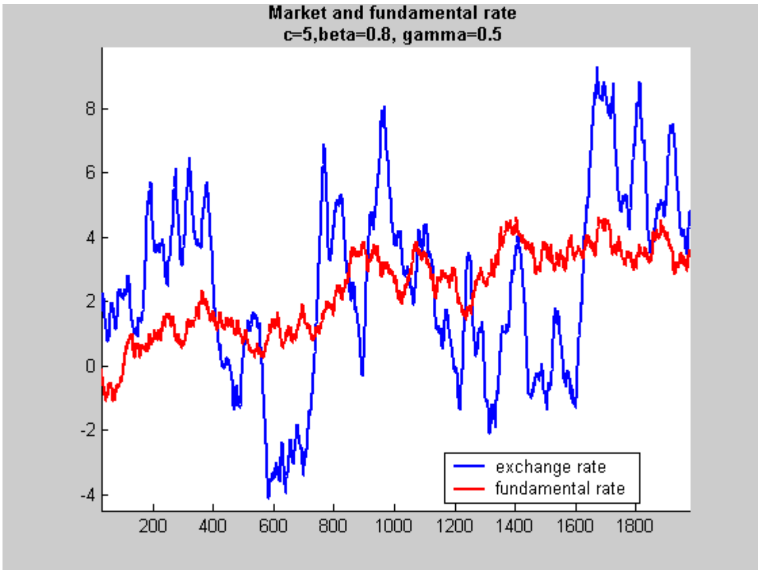

Figure 14:

fundamental using the same parameter values as in figure 14 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:

$$\Delta s_t = \mu \left(s_{t-1} - \gamma s^*_{t-1} \right) + \sum_{i=1}^{n} \lambda_i \Delta s_{t-i} + \sum_{i=1}^{n} \phi_i \Delta s^*_{t-i} \tag{16}$$

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.

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 0.3% 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 slow mean reversion of the real exchange rate is due to slow adjustment of the nominal exchange rate and not of the goods prices.

Table 1: error correction model

Error correction		$\Delta\mathbf{s}_{t-i}$				$\Delta\mathbf{s}^*_{t-i}$			
$\boldsymbol{\mu}$	$\boldsymbol{\gamma}$	$\boldsymbol{\lambda}_1$	$\boldsymbol{\lambda}_2$	$\boldsymbol{\lambda}_3$	$\boldsymbol{\lambda}_4$	$\boldsymbol{\varphi}_1$	$\boldsymbol{\varphi}_2$	$\boldsymbol{\varphi}_3$	$\boldsymbol{\varphi}_4$
-0.003	0.92	0.32	0.20	0.13	0.08	0.03	0.02	0.01	0.01
-5.9	*4.9*	*22.8*	*13.7*	*8.7*	*5.9*	*1.9*	*1.0*	*0.6*	*0.1*

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 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))[5] .

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.

6.2 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).

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:

$$var(s) = var(f) + var(n) \tag{17}$$

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 (17) we obtain

[5] 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 .

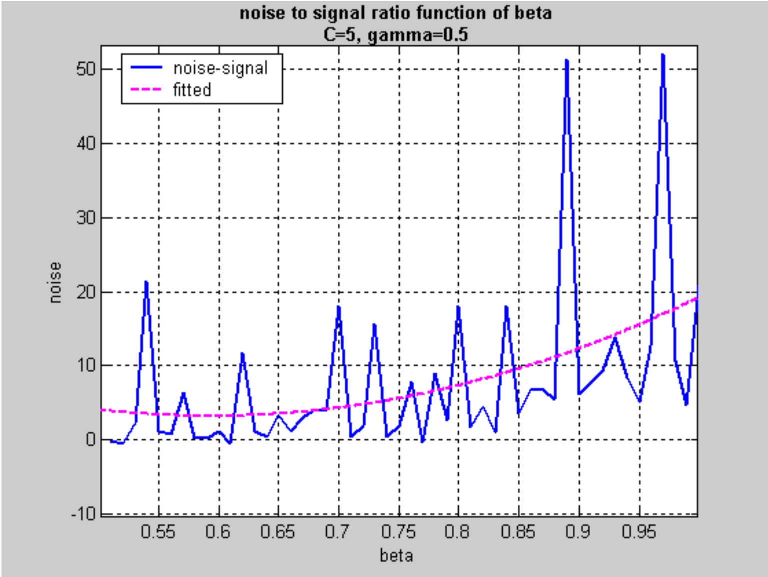

Figure 15:

$$\frac{var(n)}{var(f)} = \frac{var(s)}{var(f)} - 1 \tag{18}$$

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 15). 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 16.

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

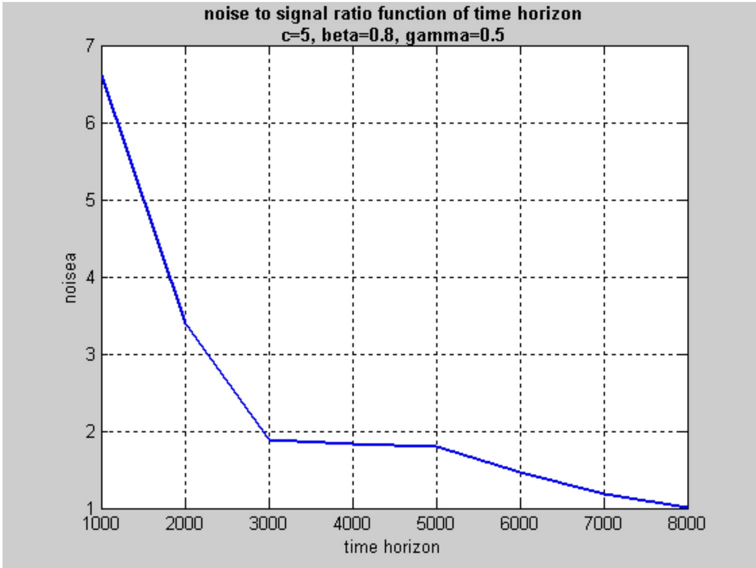

Figure 16:

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 17 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.

6.3 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

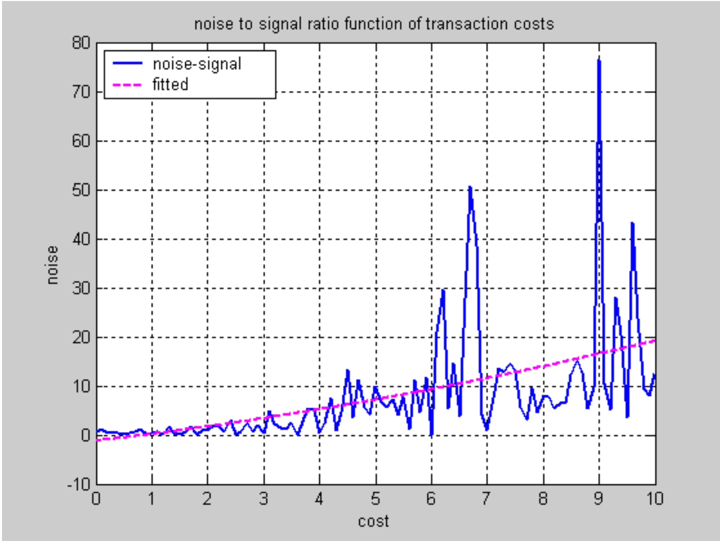

Figure 17:

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 of the kurtosis and of the Hill index in table 2. 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 4 we show the results for some sets of parameter values[6]. 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

[6] 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.

Table 2: Kurtosis and Hill index

Parameter values	kurtosis	median Hill index		
		2.5% tail	5% tail	10% tail
C=5, beta=0.9, gamma=0.5	5.65	4.92	4.98	3.98
C=5, beta=0.9, gamma=1	4.39	4.06	4.46	3.90
C=5, beta=0.9, gamma=5	6.30	4.42	3.00	2.40
C=5, beta=0.8, gamma=0.5	8.33	4.39	4.19	3.80
C=5, beta=0.8, gamma=1	7.92	4.15	4.37	3.73
C=5, beta=0.8, gamma=5	11.08	3.63	3.90	3.54

Table 3: Kurtosis and time aggregation

Parameter values	1 period returns	10 period returns	25 period returns	50 period returns
C=5, beta=0.9, gamma=0.5	5.65	5.96	3.17	3.08
C=5, beta=0.9, gamma=1	4.39	4.11	3.67	3.45
C=5, beta=0.9, gamma=5	6.30	2.77	2.15	2.19
C=5, beta=0.8, gamma=0.5	8.33	8.52	3.14	3.43
C=5, beta=0.8, gamma=1	7.92	7.39	3.28	3.30
C=5, beta=0.8, gamma=5	11.08	10.14	3.46	3.05

of the time) and turbulent periods (occurring infrequently).

6.4 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 and Teyssière(2002), Lux and Marchesi (2000)). 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 absolute value of the simulated returns and plotted them in figure 15. This figure creates the visual impression of volatility clustering. In the second step we computed the autocorrelation function (ACF) of the absolute returns of the simulated exchange rate returns for a broad range of parameter values. In figure 17 we show the ACF for our standard set of parameters. At first glance, figure 16 suggests that the ACF dies out 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 . We estimate an ARMA model on the raw returns and we find that an ARMA(2,1) performs best. The results of estimating this model on the returns are shown in appendix 1. We then tested for serial correlation in the residuals. These tests are shown in table 4. We conclude that we cannot reject serial correlation in the error term.

Moreover, we performed an ARCH test on the residuals of the simulated

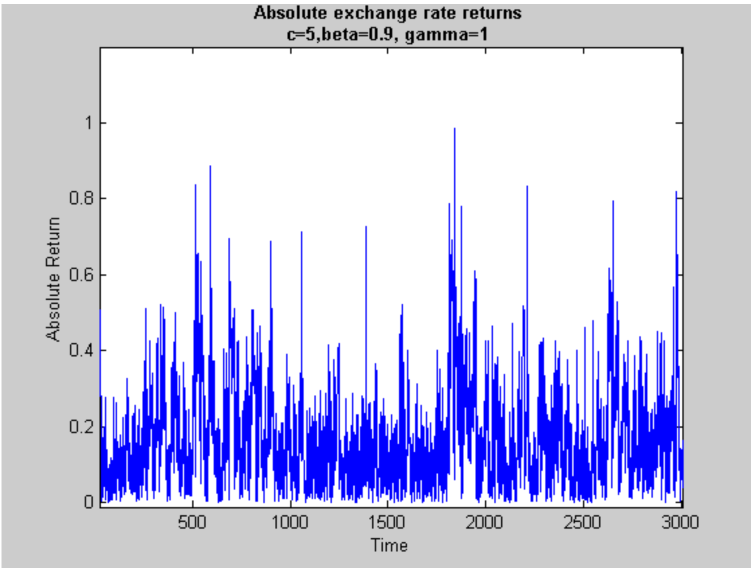

Figure 18:

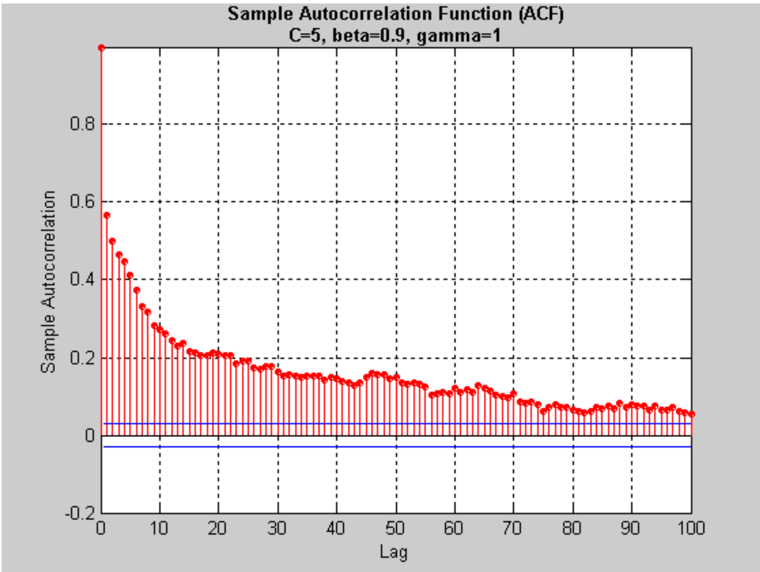

Figure 19:

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 a GARCH (2,1) specification[7]:

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

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 table **??** for different values of the extrapolation parameter β.

We observe that the GARCH coefficients, a, 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 a, δ_1 and δ_2 , which is a measure of the degree of the 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.

Table 4: Kurtosis and time aggregation

Parameter values		C=5, β=0.9, γ=1
	coefficient	T-statistic
a	0.001	0.4
b	0.001	12.4
α	0.33	13.1
δ_1	0.46	5.6
δ_2	0.12	1.9
Parameter values		C=5, β=0.9, γ=2
	coefficient	T-statistic
a	0.002	1.3
b	0.003	13.7
α	0.38	21.4
δ_1	0.23	5.2
δ_2	0.33	9.5
Parameter values		C=5, β=0.9, γ=6
	coefficient	T-statistic
a	0.001	0.2
b	0.004	10.4
α	0.35	16.9
δ_1	0.27	6.3
δ_2	0.22	6.7

[7] We also used a GARCH(1,1) specification with similar results.

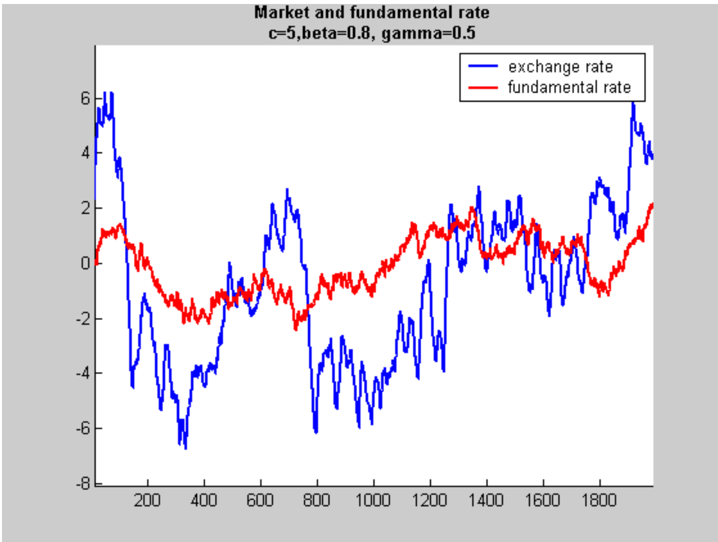

Figure 20:

7 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 20-23. (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 20 and 22). Second, the sensitivity to small changes in initial conditions is clearly visible when the variance of the exchange rate is low (see figure 21). When this variance is high, no such sensitivity can be observed (figure 23). 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 6.1 (see table 1) where this analysis refers to a low variance environment. We show the results for the high variance regime in table 6. These results contrast with those obtained in table 1. The error correction

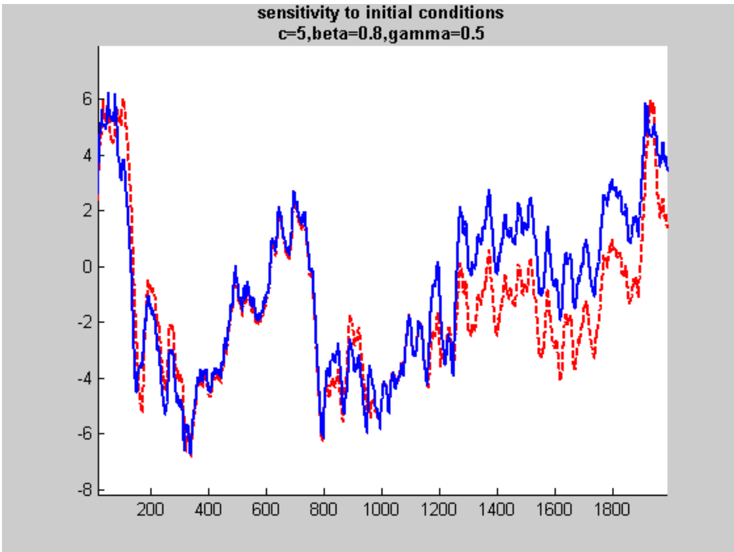

Figure 21:

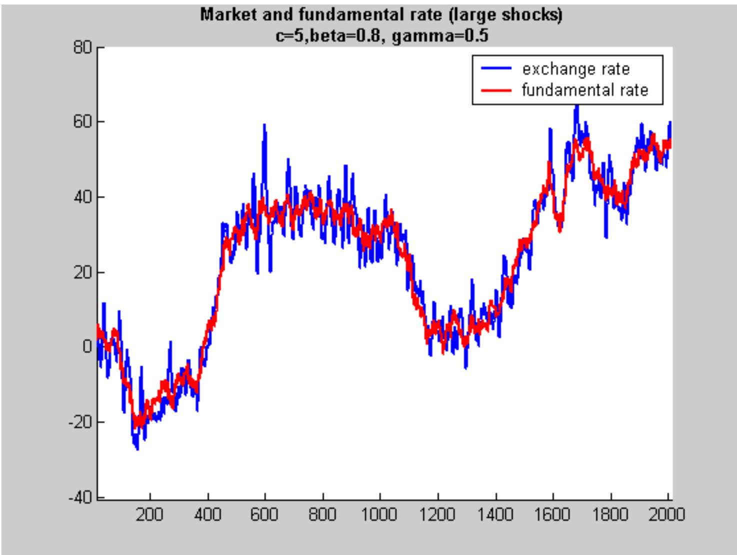

Figure 22:

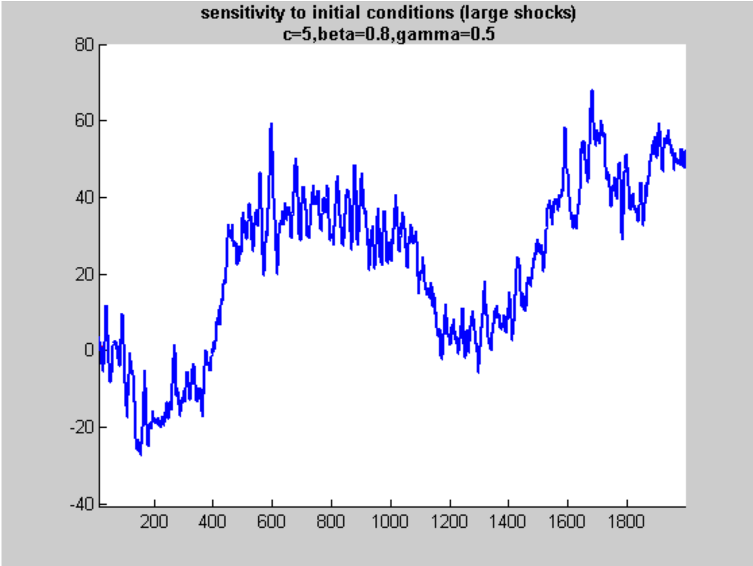

Figure 23:

coefficient is much larger in the high variance regime of table 5 than in the low variance regime of table 1. In the high variance case 9% of the deviation from equilibrium is adjusted for per period. This constrasts with a 0.3% found in the low variance case (see table 1).

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 the exchange rate and its fundamentals are always cointegrated. In the low variance case we do not find cointegration for all these sub-samples.

These results confirm what we observed from figures 18 and 19, i.e. that in a regime of high variance of shocks the exchange rate is more tightly linked to the fundamentals, and that the speed of adjustment towards the equilibrium is higher than in low variance regimes.

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 prevaling 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.

Table 5: error correction model high variance

Error correction		Δs_{t-i}				Δs^*_{t-i}			
μ	γ	λ_1	λ_2	λ_3	λ_4	φ_1	φ_2	φ_3	φ_4
-0.09	1.002	0.36	0.19	0.09	0.07	0.05	0.01	0.00	-0.01
-24.2	*156.2*	*27.5*	*13.2*	*6.7*	*5.1*	*3.3*	*0.9*	*0.06*	*-0.8*

It is also interesting to analyse how the variance regimes affect the noise-to-signal ratio. We show the results in figures 24 and 25 where we plot the noise-to-signal ratio against the extrapolation parameter β.

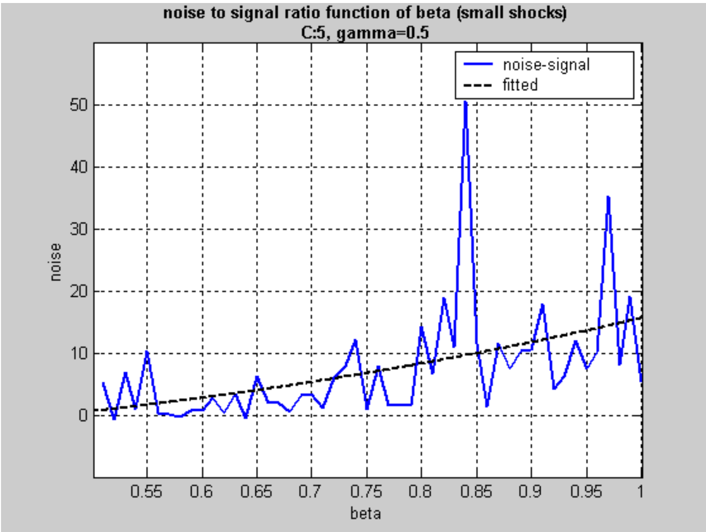

Figure 24:

In a low variance environment the noise-to-signal ratio is very large exceeding one by a wide margin. In contrast, in a high variance environment the noise-to-signal ratio is below one for most values of β. This confirmes our previous 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.

8 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

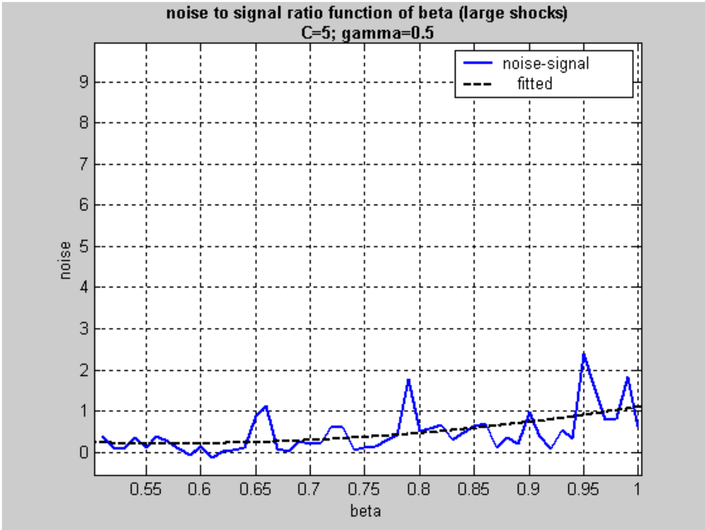

Figure 25:

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 figure 26 we show the average chartists' weight for increasing values of the intensity of choice parameter g in two different environments concerning the variance of the shocks in fundamentals. We obtained the chartists weights by simulating the model over 10000 periods and 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. On the contrary, we find that the chartists weight fluctuates around a market share, which exceeds 50%. These results are consistent with the empirical evidence of the importance of chartism in foreign exchange market (Taylor and Allen (1992)). Second, the market share of chartists increases with the parameter γ which measures the intensity with which agents react to relative profitability of the two rules. This result is related to the fact that increases in γ increase the volatility of the exchange rate. The increase in volatility is what chartists thrive on. Third, we find that, in general, chartism is more profitable in the low variance environment than in the high variance environment. This due to the fact that in a low variance environment the exchange rate movements are disconnected from the fundamental most of the time making fundamentalist forecasting relatively unattractive. We show this feature in figure 27 which presents the average profits of chartists and fundamentalists in the low variance environment. It is

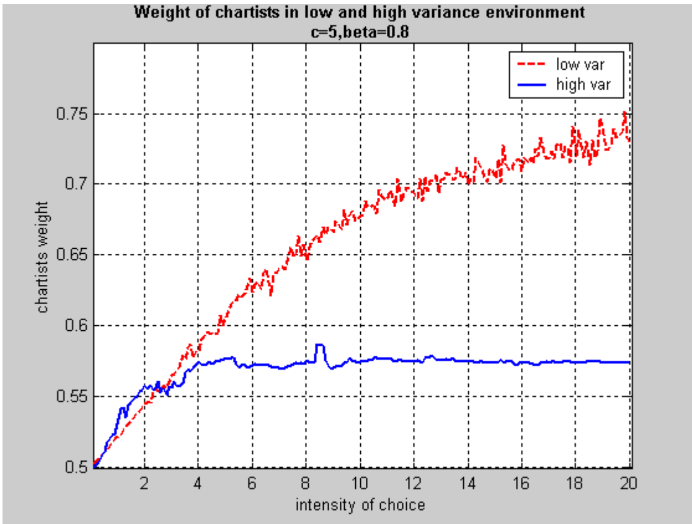

Figure 26:

striking to see that in such an environment chartism appears to be much more profitable than fundamentalism.

From the preceding analysis we conclude that chartism is evolutionary stable, and that it is generally more profitable than fundamentalism . In addition, there is a positive correlation between turbulence (noise) and the share of chartism in the market. With more noise there is more chartist profit and thus more chartists. The reverse is also true: with more chartists there is more noise and thus more profits for chartists. These results suggests that there is a selffulfilling evolutionary dynamics present in the system which can be described as follows. As the chartists increase in numbers, the noise they create makes the use of chartists rules more profitable. At the same time, the chartists have the effect of "creating smoke around the fundamentals", making fundamentalists' forecasting less profitable. 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 chartists drive 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. 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 prevent the chartists from completely

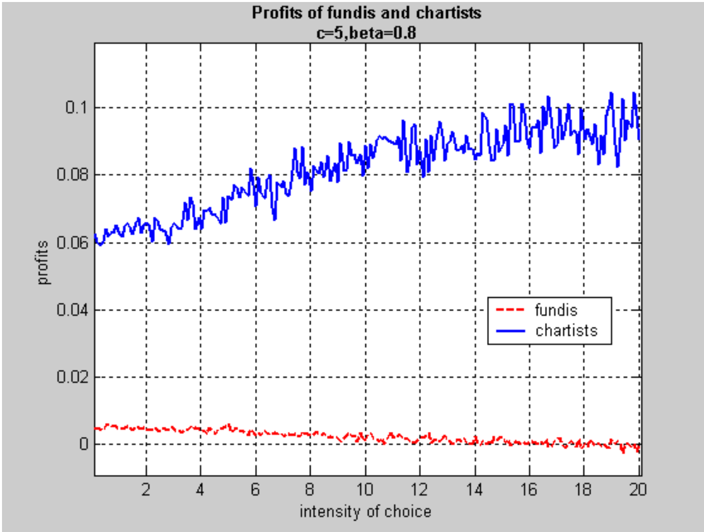

Figure 27:

driving out the fundamentalists and taking over the market.

9 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 in the goods markets produce an important non-linearity in the model. Agents are heterogeneous in that they have different beliefs, i.e. they use different forecasting rules. The relative importance of these different types of agents is driven by the relative profitability of their forecasting rules and by the risk associated with the use of such rules. Thus agents are rational in the sense that they evaluate ex post the relative (risk-adjusted) profitability of the forecasting rules and switch to the better one. We argued that this trial and error process is a better way to model agents' behaviour than to assume that their expectations are rational. In rational expectations models with heterogeneous agents the burden of collecting and processing information for individual agents is extraordinarily high. Not only must individual agents know the structure of the model, but they must also be able to read the minds of all the other agents. In traditional religions agents with such intellectual capacities were called Gods.

Our 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 contrasts 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. We argued that the traditional rational expectations efficient market model has failed empirically. 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 mimics 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 tranquillity and turbulence is itself unpredictable.

>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-fulfilling 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.

10 References

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

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., 1997, A Rational Route to Randomness, Econometrica, 65, 1059-1095

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.

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

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.

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.

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.

Johansen, S., 1991, Estimation and Hypothesis Testing of Cointegration Vectors in Gaussian Vector Autoregressive Models, Econometrica, 55, 1551-80.

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.

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.

11 APPENDIX 1: The ARMA model of the exchange rate returns

Table 6: ARMA model

variable	**coefficient**	**t-statistic**
constant	-0.003	-0.28
AR(1)	0.81	19.96
AR(2)	0.09	2.72
MA(1)	-0.33	-8.39
R^2	*0.60*	