

Real Exchange Rates and Switching Regimes

U. MICHAEL BERGMAN*

JESPER HANSSON[†]

Revised, June 2000

Abstract

We suggest that the real exchange rate between the major currencies in the post-Bretton Woods period can be described by a stationary, two state Markov switching AR(1) model. Based on the forecast performance we find that this model out-performs two competing models where the real exchange rate is non-stationary. We also find that the existence of different regimes, as in the Markov switching model, is consistent with the common finding of unit roots in real exchange rates.

KEYWORDS: Real exchange rates; Markov switching autoregressive models; forecasts, simulation

JEL CLASSIFICATION: C22, C53, F31

1 Introduction

The view that the law of one price, or its aggregate equivalent purchasing power parity, holds in the long-run but not necessarily in the short-run is wide spread among economists.¹ This view is, however, not supported by strong empirical evidence at least

*Department of Economics, Lund University, P.O. Box 7082, S-22007 Lund, Sweden.

[†]Department of Economics, Lund University, P.O. Box 7082, S-22007 Lund, Sweden.

This paper was largely written while Michael Bergman visited Economic Policy Research Unit (EPRU) at the Copenhagen Business School. He is grateful to the institute for its hospitality. We have received valuable comments from Menzie Chinn, Kon Lai, two anonymous referees and seminar participants at EPRU, Lund University and University of California, Santa Cruz. We would also like to thank Anders Warne for generously supplying his Gauss code for the EM algorithm.

¹See Froot and Rogoff (1995) or Rogoff (1996) for recent evaluations of the purchasing power parity doctrine and the empirical evidence.

when examining the most recent period with floating exchange rates.² Usually, it is difficult to reject the null hypothesis that the real exchange rate contains a unit root. A potential reason for not rejecting unit roots is that standard tests usually have low power, in particular for near unit root processes.

The recent literature has suggested several solutions to this low power problem. For example, Perron (1990) has shown that standard unit root tests are biased towards a non-rejection of the null in the presence of a change in the intercept.³ In more recent papers, Perron and Vogelsang (1992) and Dropy (1996) develop unit root tests allowing for one-time changes in the mean. When applying their test statistics to real exchange rates, they find that unit roots can indeed be rejected when allowing for structural instability. Empirical evidence suggesting structural instability is also found when comparing out-of-sample forecasts based on different estimation sample lengths, see Siddique and Sweeney (1998).

From an empirical point of view, it seems very restrictive to assume that there is only a one-time break in the intercept. It may actually be the case that the intercept varies over time where the different intercepts are associated with different regimes or states of nature. Another possibility is that all parameters of the model, e.g. the intercept and the variance, change over time. In the context of nominal exchange rates, it is natural to distinguish between two phases, exchange rate appreciation and depreciation regimes. To model such long-swings in the data, Engel and Hamilton (1990) suggest a Markov switching random walk model with drift. They examined the out-of-sample forecast behavior of this model and compared with a single regime random walk model. For the three nominal exchange rates in their sample, they found that mean squared errors from the Markov switching model in many cases were lower than for the single regime random walk model.⁴

In this paper we suggest that the level of the real exchange rate is generated by a stationary 2-state Markov switching autoregressive model of the first order, a model nested within the class of models studied by Engel and Hamilton (1990) and Engel (1994). We extend the earlier literature focusing on one-time structural shifts in the mean by allowing all parameters (the mean, the autoregressive parameter and the variance) to vary

²See for example, Cheung and Lai (1993a, 1993b), Juselius (1995), Ott (1996) and Lothian (1997).

³Other approaches to solve low power problems when testing for PPP include using simulations (e.g. Edison, Gagnon and Melick (1997)), multivariate maximum likelihood methods (e.g. Cheung and Lai (1993b)), multivariate versions of standard unit root tests (e.g. Jorion and Sweeney (1996) and Lothian (1997)), panel data regressions (e.g. Oh (1996)) and non-linear models (e.g. Creedy, Lye and Martin (1996) and Sarantis (1999)). However, Engel (2000) shows that tests of PPP have important size biases and argues that PPP may not hold after all.

⁴However, later research by Engel (1994) showed that the single regime random walk model often outperformed the Markov model for other nominal exchange rates.

across different states where the process driving the state is a Markov switching process.⁵ The performance of this base line model is evaluated by comparing mean squared errors of out-of-sample forecasts at different forecast horizons with both a single regime and a 2-state Markov switching random walk model.⁶

Our assumption that the real exchange rate is generated by a 2-state Markov switching AR(1) model is consistent with predictions of the general equilibrium model examined by Dumas (1992). Under the assumption of spatially separated countries and shipping costs it is shown that the real exchange rate switches between two states and exhibits mean reversion within each regime where the drift is negative (positive) when the deviation from parity is positive (negative). Dumas (1992) also show that the deviation from parity is heteroscedastic. These theoretical predictions imply certain restrictions on our Markov switching model, the intercept, the autoregressive parameter and the innovation variance should all depend on the state.

The empirical results suggest that we can reject a single regime autoregressive model for the real exchange rates and that only the drift is state dependent. When comparing in-sample as well as post-sample forecasts, we find that the 2-state Markov switching model out-perform forecasts from both types of random walk models. In addition to the forecasting comparison, we also examine whether our stationary 2-state Markov switching AR(1) model is capable of explaining results obtained using standard unit root tests such as the augmented Dickey-Fuller (ADF) test. For this experiment we assume that our base model is the “true” data generating process such that we can generate artificial real exchange rate data. Then we test whether simulated data contain unit roots applying the standard ADF-test. Similar to the results in Perron (1990), we find that the probability of rejecting a unit root is very low. This suggests that the ADF-test is biased towards a non-rejection of the null when the data is generated by a stationary Markov switching AR(1) model. It is therefore not surprising that standard tests tend to reject unit roots if real exchange rates are generated by a stationary Markov switching AR(1) model.

The remainder of the paper is structured in the following manner. In section 2 we introduce the Markov switching model. Section 3 presents the empirical analysis. Finally, section 4 concludes the paper.

⁵In a recent paper, Engel and Kim (1999) examine the behavior of the pound/dollar real exchange rate under the assumption that the real exchange rate is integrated and assuming that the deviation from the permanent component can be modelled as a 3-state Markov switching model.

⁶Creedy, Lye and Martin (1996), Norrbin, Reffett and Ji (1997) and Fritsche and Wallace (1997) also compare forecast behavior of different versions and extensions of PPP. The use of forecasting behavior as way of distinguishing between different model specifications in the context of exchange rates was initiated by Meese and Rogoff (1983).

2 The Markov switching AR(1) model

In this section we present a univariate first order autoregressive, AR(1), model subject to changes in regime. The difference between our model and the model used by Engel and Hamilton (1990) and Engel (1994) is that we allow the autoregressive parameter to be different from unity. The parametric restrictions we impose imply ergodicity of the regime process and weak stationarity of the observable variable. Let x_t denote the logarithm of the real exchange rate, i.e., $x_t = e_t + p_t^* - p_t$, where e_t is the logarithm of the nominal exchange rate, p_t^* is the logarithm of the foreign price level, p_t is the logarithm of the domestic price level. Assume then that the time series x_t is generated by the following Markov switching AR(1) model:

$$x_t = \mu_{s_t} + \phi_{s_t} x_{t-1} + \varepsilon_t, \quad (1)$$

where $\varepsilon_t \mid s_t \sim N(0, \sigma_{s_t}^2)$ and the initial value x_0 is taken as fixed. The unobserved random state variable s_t is independent of past x_t conditional on s_{t-1} and takes on the value 1 or 2 and is assumed to follow a 2-state Markov process, i.e.,

$$\Pr\{s_t = j \mid s_{t-1} = i\} = p_{ij} \quad \text{for } i \text{ and } j = 1, 2. \quad (2)$$

The Markov transition probabilities satisfy $\sum_{j=1}^2 p_{ij} = 1$ for $i = 1, 2$. These transition probabilities, assuming no absorbing states, can be collected in the 2×2 transition matrix P ,

$$P = \begin{bmatrix} p_{11} & 1 - p_{22} \\ 1 - p_{11} & p_{22} \end{bmatrix}. \quad (3)$$

A sufficient condition for x_t to be weakly stationary, given by Karlsen (1990), is that the 2×2 matrix

$$A = \begin{bmatrix} \phi_1^2 p_{11} & \phi_1^2 p_{21} \\ \phi_2^2 p_{12} & \phi_2^2 p_{22} \end{bmatrix}$$

has all eigenvalues inside the unit circle. Note that this stationarity condition implies that it is the interaction between the autoregressive parameters and the transition probabilities that determines whether the variable x_t is weakly stationary or not. It may well be the case that any of the two autoregressive parameters is greater than one and still the stationarity condition above is satisfied, see Holst, Lindgren, Holst and Thuvsholmen (1994). The intuition is that as long as there is a sufficiently high probability of switching from the explosive regime to a stationary regime, the whole model is stationary.

Note that the Markov chain can be represented as a vector autoregressive process by letting ξ_t denote a random 2×1 vector defined as

$$\xi_t = \begin{cases} \begin{bmatrix} 1 & 0 \end{bmatrix}' & \text{when } s_t = 1 \\ \begin{bmatrix} 0 & 1 \end{bmatrix}' & \text{when } s_t = 2 \end{cases}. \quad (4)$$

If $s_t = 1$, then the second element of ξ_{t+1} is a random variable that takes on the value 1 with probability $p_{12} = 1 - p_{11}$ and zero otherwise. The conditional expectation of ξ_{t+1} given $s_t = 1$ is therefore given by the first column of the transition matrix P , i.e.,

$$\mathbb{E}[\xi_{t+1} | s_t = 1] = \begin{bmatrix} p_{11} \\ 1 - p_{11} \end{bmatrix}. \quad (5)$$

From (4) and (5) using the Markov property described in (2) it follows that

$$\mathbb{E}[\xi_{t+1} | \xi_t, \xi_{t-1}, \xi_{t-2}, \dots] = P\xi_t, \quad (6)$$

which further implies that it is possible to represent the Markov chain as the first order vector autoregression

$$\xi_{t+1} = P\xi_t + v_{t+1}, \quad (7)$$

where $v_{t+1} \equiv \xi_{t+1} - \mathbb{E}[\xi_{t+1} | \xi_t, \xi_{t-1}, \dots]$ is a martingale difference sequence.

Conditional on knowing s_{t+1} , our model in (1) can be used to compute one step ahead forecasts of x_t in the standard way,

$$\mathbb{E}[x_{t+1} | s_{t+1} = j, Y_t; \theta] = \mu_j + \phi_j x_t, \quad (8)$$

where $Y_t = [x_0, \dots, x_t]'$ denotes the information set as of date t and

$$\theta = [\mu_1 \quad \mu_2 \quad \phi_1 \quad \phi_2 \quad \sigma_1^2 \quad \sigma_2^2 \quad p_{11} \quad p_{22}]'$$

is a vector containing the parameters of the Markov switching model. In order to compute the forecast, we also need a forecast of which regime is likely to govern the process in the future. From (7) we have that the state variable is governed by

$$\xi_{t+m} = v_{t+m} + Pv_{t+m-1} + P^2v_{t+m-2} + \dots + P^{m-1}v_{t+1} + P^m\xi_t, \quad (9)$$

implying that the optimal m -period-ahead forecast of ξ_{t+m} is given by

$$\mathbb{E}[\xi_{t+m} | Y_t; \theta] = P^m \mathbb{E}[\xi_t | Y_t; \theta]. \quad (10)$$

Note that even though we use all information available at date t , the expected value of ξ_t is an estimate since the state variable s_t is unobservable. This implies that the expected value of ξ_t can be interpreted as the conditional probability of being in regime one and two respectively, i.e.,

$$\mathbb{E}[\xi_t | Y_t; \theta] = \begin{bmatrix} \Pr\{s_t = 1 | Y_t; \theta\} \\ \Pr\{s_t = 2 | Y_t; \theta\} \end{bmatrix}. \quad (11)$$

To obtain the forecast of x_{t+1} we just multiply the appropriate forecast for the j th regime given by equation (8) with the probability that the process will be in that regime

given by equation (10) and sum those products for every regime, i.e.,

$$\begin{aligned} E[x_{t+1} | Y_t; \theta] &= \sum_{j=1}^2 E[x_{t+1} | s_{t+1} = j, Y_t; \theta] \Pr\{s_{t+1} = j | Y_t; \theta\} \\ &= h_{t+1}' E[\xi_{t+1} | Y_t; \theta] = h_{t+1}' P^m E[\xi_t | Y_t; \theta], \end{aligned}$$

where h_{t+1} is a 2×1 vector containing the regime specific conditional forecasts given by (8). Multiperiod forecasts can then be computed by recursive substitution using the formula for the one-step-ahead forecasts. The two periods ahead forecast is, for instance, given by

$$\begin{aligned} E[x_{t+2} | Y_t; \theta] &= \sum_{j=1}^2 E[x_{t+2} | s_{t+2} = j, Y_t; \theta] \Pr\{s_{t+2} = j | Y_t; \theta\} \\ &= \begin{bmatrix} \mu_1 + \phi_1 h_t' E[\xi_{t+1} | Y_t; \theta] \\ \mu_2 + \phi_2 h_t' E[\xi_{t+1} | Y_t; \theta] \end{bmatrix}' P^2 E[\xi_t | Y_t; \theta]. \end{aligned}$$

By using these forecasts as a tool for evaluation of the model, cf. Meese and Rogoff (1983), we can evaluate and compare out-of-sample forecasts from the Markov switching model with that of alternative models. In particular, we compare forecasts from our Markov model to forecasts from a single regime random walk with drift model, whose forecasts are given by $E[x_{t+m} | Y_t; \mu_{\Delta x}] = x_t + m\mu_{\Delta x}$ where $\mu_{\Delta x}$ is replaced by the estimated mean $\sum_{t=1}^T \Delta x_t / T$, and a 2-state Markov switching random walk model where the autoregressive parameter $\phi_i = \phi = 1$.⁷ For both these alternative models, the real exchange rate is a nonstationary process.

3 Empirical Work

3.1 Data

The data consists of quarterly observations on the period average spot exchange rates (in units of foreign currency per US dollar) and the consumer price index for six major industrialized countries (the UK, France, Germany, Switzerland, Canada, Japan and USA) taken from the IMF International Financial Statistics CD-ROM. The sample runs from the second quarter of 1973 to the fourth quarter of 1997 using the first quarter of 1973 as initial condition.⁸ The effective sample of observations is 99. The real exchange rate

⁷These alternative models have been used by Engel and Hamilton (1990) and Engel (1994) to model nominal exchange rates.

⁸The reason why we use period averages instead of end-of-period nominal exchange rate is that the estimated models always pass diagnostic testing. This is not always the case when using end-of-period data. However, the main conclusions from our empirical analysis are unaffected.

is normalized to unity in the second quarter of 1973 and we use 100 times the natural logarithm of the real exchange rate in our empirical analysis.

The logarithm of the six real exchange rates is depicted in Figure 1. The four European real exchange rates show a great deal of variability over time with a declining trend until 1979/80 and then a sharp rising trend which was broken in 1985. The real exchange rates then fell sharply the next few years and then they declined more moderately through 1997. It is possible that these similarities reflect the effects of European currency cooperation and the increased integration of European economies. Long swings are also apparent for the Canadian dollar and the Japanese Yen real exchange rates but in general the behavior of these series are quite different from the European ones. However, the rising trend in all six real exchange rates was broken in 1985. It is also notable that the long swings in nominal exchange rates documented in, for example, Engel and Hamilton (1990) also show up in real exchange rates. This suggests that relative prices are fairly stable over time which would indicate that the variability in nominal exchange rates is not induced by corresponding changes in relative prices as suggested by the PPP theory.

3.2 Estimates of the Markov switching AR(1) model

The first step in our empirical analysis is to specify a 2-state Markov switching AR(1) model for the six real exchange rates, British Pound, French Franc, German Mark, Swiss Franc, Canadian Dollar and the Japanese Yen. We estimate each model independently, thus allowing for different switches between the two states across the six real exchange rates. All models are estimated using the first 72 observations, i.e., the effective sample runs from 1973:2 to 1990:4. We keep the remaining 28 observations of our sample (1991:1–1997:4) for the evaluation of post sample forecasts.

Table 1 shows maximum likelihood estimates of our base model for the six real exchange rates allowing only the intercept to vary across the two states.⁹ The point estimates suggest that the first state is associated with a rising real exchange (real depreciation) rate whereas the second state is associated with a declining real exchange rate (real appreciation), the estimated intercept is positive for state one whereas it is negative for state two, except for the Japanese Yen where the intercept in state one is not significantly different from zero. Both the estimate of the autoregressive parameter and the eigenvalue

⁹Initially, we allow all parameters, the intercept, the autoregressive parameter and the variance to vary across the two states. Within this model, we then test whether the parameters do not vary across the states. The maintained assumption is that there are two unobserved states which are assumed to follow a 2-state Markov chain. Wald and likelihood ratio tests of both the individual and joint null hypotheses that the autoregressive parameter and the variance are state independent cannot be rejected at the ten percent significance level for any real exchange rate except for the French Franc where the Wald test, but not the LR-test, rejects the null hypothesis that the variance is state independent. These results are available from the authors on request.

Figure 1: Real exchange rates 1973:2–1997:4.

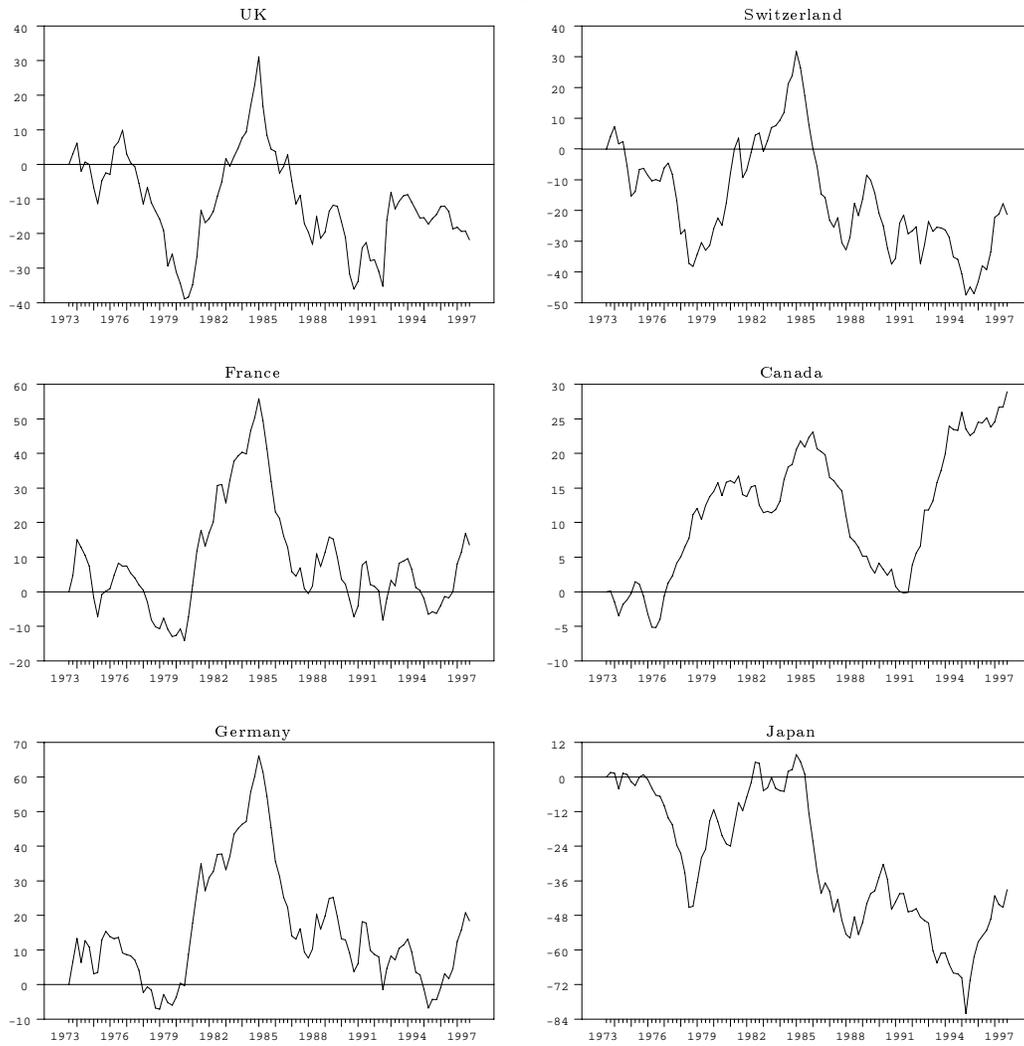


Table 1: Estimates of a 2–state Markov switching AR(1) model for six real exchange rates 1973:2–1990:4.

Parameter	UK	France	Germany	Switzerland	Canada	Japan
μ_1	3.554 (0.550)	6.131 (0.604)	6.569 (0.733)	2.390 (0.726)	1.693 (0.230)	−0.370 (0.681)
μ_2	−5.096 (0.549)	−2.845 (0.409)	−2.676 (0.487)	−6.556 (0.775)	−0.306 (0.249)	−8.932 (1.157)
ϕ	0.928 (0.027)	0.904 (0.020)	0.888 (0.023)	0.958 (0.027)	0.922 (0.021)	0.871 (0.027)
σ^2	10.118 (1.698)	7.706 (1.293)	10.719 (1.799)	13.513 (2.268)	1.644 (0.276)	15.879 (2.665)
p_{11}	0.672	0.679	0.682	0.792	0.952	0.911
p_{22}	0.690	0.833	0.830	0.716	0.944	0.817
$\max \text{eig}(\hat{A}) $	0.861	0.817	0.789	0.918	0.850	0.759
Autocorr.	[0.632]	[0.729]	[0.281]	[0.260]	[0.413]	[0.859]
ARCH	[0.983]	[0.251]	[0.564]	[0.689]	[0.237]	[0.263]
Markov	[0.542]	[0.165]	[0.486]	[0.937]	[0.383]	[0.497]

Note: Standard errors in parentheses and p-values in square brackets.

statistic suggest that the six real exchange models are stationary. This is an interesting finding as it suggests that the level of real exchange rates is stationary when allowing for multiple shifts in the intercept. Moreover, our finding is consistent with the empirical evidence in Perron and Vogelsang (1992) who found evidence suggesting stationarity around a one–time shift in the intercept.

To evaluate whether our Markov model is correctly specified, we test for autocorrelation and ARCH in the residuals. These tests are based on the conditional scores following the suggestions made by Hamilton (1996). We also test the so called Markov assumption, i.e., the hypothesis that the state variable s_t only display first order serial correlation as described by (7). The lower panel of Table 1 reports p-values from tests of the null hypothesis of no autocorrelation, no ARCH and the Markov assumption. We note that these hypotheses cannot be rejected at conventional significance levels for any country.

To formally test the null hypothesis that the six real exchange rates are generated by a one–state AR(1) model against a 2–state Markov switching model where only the intercept is state dependent, we apply the test procedure suggested by Hansen (1992, 1996). Since the transition probabilities are not identified under this null hypothesis, Hansen suggests that a grid search should be used where the likelihood function is evaluated for different

values of the transition probabilities (p_{11} and p_{22}) and the switching parameter (the drift μ). We use the following three grids for the transition probabilities:

Grid 1 for p_{11} and p_{22} : 0.20 to 0.970 in steps of 0.11 (eight gridpoints)

Grid 2 for p_{11} and p_{22} : 0.15 to 0.975 in steps of 0.075 (12 gridpoints)

Grid 3 for p_{11} and p_{22} : 0.18 to 0.980 in steps of 0.05 (17 gridpoints)

These grids include the maximum likelihood estimates of the transition probabilities for the six real exchange rates reported in Table 1. For the drift μ , we use the following grids with 30 gridpoints:

Grid for United Kingdom: -7.0 to 4.6 in steps of 0.4

Grid for France and Germany: -4.0 to 7.6 in steps of 0.4

Grid for Switzerland: -8.0 to 3.6 in steps of 0.4

Grid for Canada: -2.0 to 3.8 in steps of 0.2

Grid for Japan: -10.0 to 1.6 in steps of 0.4

In Table 2 we report the LR test statistics of the null hypothesis that the real exchange rates are generated by a one-state AR(1) model against the alternative of a 2-state Markov switching model where only the intercept is state dependent. From this table we note that we can reject the null hypothesis for the French Franc, German Mark, Swiss Franc, Canadian dollar and the Japanese Yen at the 1 percent level. For the British Pound, we reject the null at the 11 percent level. Based on these estimates, we conclude that it is unlikely that the real exchange rates (with the possible exception of the British Pound) are generated by a one-state AR(1) model.

From the maximum likelihood estimates, using the full sample of observations for each real exchange rate, we can draw inference concerning the state process at each time, i.e., $\Pr\{s_t = 1|Y_T; \theta\}$.¹⁰ Figure 2 plots the six real exchange rates and these smoothed probabilities. Shaded vertical bars in this figure show that the system is in state one based on $\Pr\{s_t = 1|Y_T; \theta\} > 0.5$. A striking feature of these graphs is that our base line model seems to capture the long swings that are apparent in the data. As the maximum

¹⁰All parameter estimates, using the full sample 1973:2-1997:4, are very similar to the estimates for the sub sample 1973:2-1990:4 shown in Table 1, except for the British Pound where, in particular, the transition probability p_{22} is considerably higher, 0.903 for the full sample compared to 0.690 for the sub sample. The intercepts are also larger in absolute values in both regimes for the full sample. The maximum likelihood estimates for the other five real exchange rates are almost identical to the parameter estimates in Table 1.

Table 2: Likelihood ratio tests of the null hypothesis that the real exchange rate follows a one-state AR(1) model.

	UK	France	Germany	Switzerland	Canada	Japan
Grid 1	2.780 (0.101)	4.732 (0.000)	4.066 (0.001)	4.063 (0.001)	4.224 (0.001)	4.383 (0.001)
Grid 2	2.847 (0.096)	4.809 (0.001)	4.182 (0.001)	4.138 (0.001)	4.282 (0.002)	4.442 (0.002)
Grid 3	2.780 (0.105)	5.965 (0.000)	4.658 (0.000)	4.998 (0.000)	4.224 (0.001)	5.305 (0.000)

Note: P-values are shown in parentheses below each test statistic.

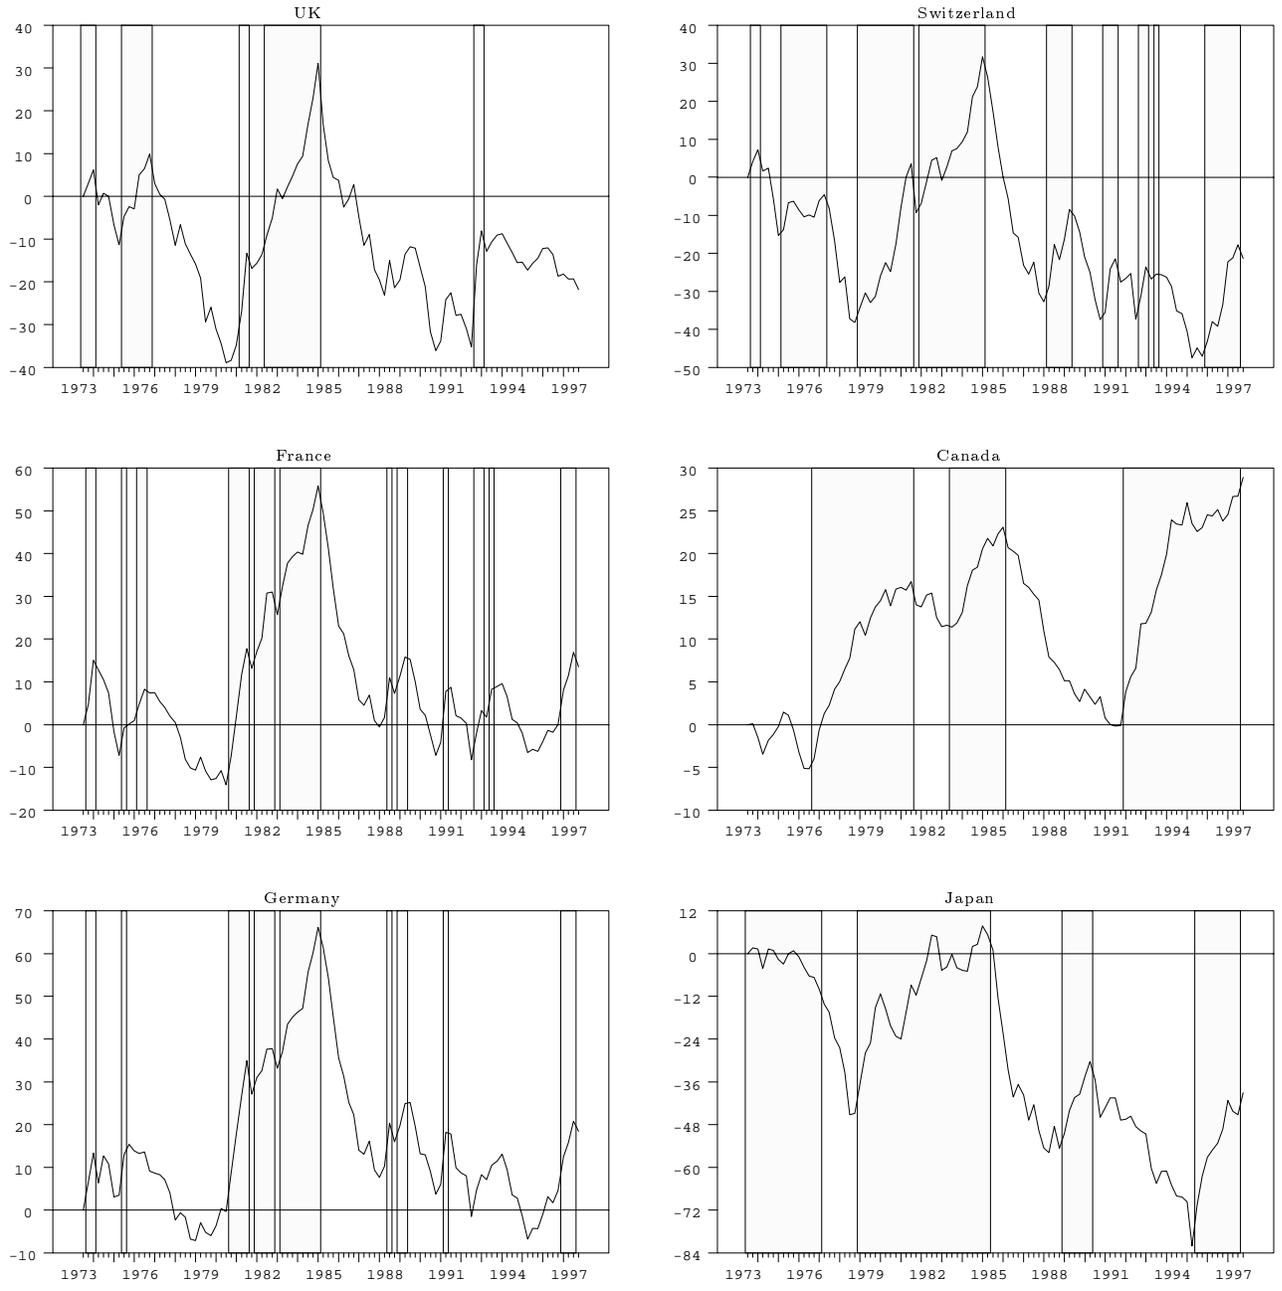
likelihood estimates in Table 1 suggest, state one is associated with rising real exchange rates whereas state two is associated with declining real exchange rates. Figure 2 also reveals that regime switches appear quite frequently for three real exchange rates, the German Mark, the French Franc and the Swiss Franc. This reflects the estimates of the transition probabilities in Table 1, where the probability that the system remains in state one or state two is relatively low, between 0.67 and 0.83. For the other three real exchange rates, there is a higher probability that the system stays in the same state which implies fewer switches of the regime.

3.3 Forecasts of the model

In this subsection we examine the forecasting performance of our base line model, the 2-state Markov switching AR(1) model (MS-AR(1)) and compare these forecasts with forecasts from both a single regime random walk with drift model (RW) and a 2-state Markov switching random walk with drift model (MS-RW). Engel and Hamilton (1990) used the latter model to study nominal exchange rates and it is interesting to investigate the forecasting performance of this model as well. The model is identical to the MS-AR(1) model with the autoregressive parameter constrained to unity and imposing the restriction of equal variances across the two states.¹¹ If the forecasts from the stationary Markov switching AR(1) model have lower mean squared errors than the forecasts from the other two models, we suggest that this model is superior to the two alternative models. The post-sample forecasts are based on the Maximum likelihood estimates in Table 1.

¹¹We cannot reject the null hypothesis that the variance is constant across the two states using Wald and LR tests at the 5 percent level for any real exchange rate. These results are available on request from the authors.

Figure 2: Probability that $s_t = 1$ conditional on all information contained in the sample period and the actual real exchange rate. Shaded bars indicate $\Pr \{s_t = 1 | Y_T; \theta\} > 0.5$.



Using these estimates, we construct the regime probabilities $\Pr\{s_t = j|Y_t; \theta_0\}$ and then we compute mean squared forecast errors for the post-sample period 1991:1 – 1997:4.

In Table 3 we compare the out-of-sample forecast performance with the RW and MS-RW models at four different forecast horizons, one to four quarters ahead.¹² The first row for each country present mean squared forecast error from a single regime RW model, while the following two rows show improvements in percent of the forecasts from the competing models.

In most cases, see Table 3, forecasts from the MS-AR(1) model are superior to both random walk models with improvements ranging from four to 12 percent at the one quarter horizon and from two to 23 percent at the four quarter horizon. The improvements of our base line model over the random walk model also tend to increase with the forecast horizon. The forecast performance of the MS-RW model, on the other hand, tends to deteriorates with the forecast horizon in most cases. The only exceptions are the Canadian dollar, where the improvements of forecasts over the single random walk model are of about the same magnitude for both Markov switching models, and the Japanese Yen where forecasts generated by the MS-AR(1) model deteriorates with the forecast horizon, see Table 3.

To evaluate whether the improvements in forecasting performance reported in Table 3 are statistically significant, we apply the modified Diebold-Mariano test for forecast encompassing suggested by Harvey, Leybourne and Newbold (1998). Forming a composite forecast as a weighted average of, for example, the forecasts using the MS-AR(1) model and the RW model, the preferred forecast encompasses the competing forecast if the optimal weight on the latter is zero. In Table 4, we report modified Diebold-Mariano test statistics for forecast encompassing. The first two rows for each country contain tests comparing forecasts using the MS-RW model and forecasts using the single regime RW model. In general, the results do not indicate which model produces better forecasts. We reject the null hypothesis that the RW model encompasses the MS-RW model in five out of 24 cases, using the 10 percent significance level. When switching the null hypothesis, we reject the null that the MS-RW model encompasses the RW model in five out of 24 cases. In the remaining 14 cases, neither the RW model nor the MS-RW model encompasses the alternative model.

When comparing forecasts using our baseline model (MS-AR(1)) and the RW model, we find that the former encompasses the latter model for all countries except for Japan. The alternative hypothesis that the RW model encompasses the MS-AR(1) model is rejected in 17 out of 24 cases at the 10 percent significance level and in all cases the p-values are less than 13 percent level, see Table 4. Note also that the tests of encompassing for the Japanese Yen are inconclusive. We reject the null that the MS-AR(1) model encom-

¹²We have also computed in-sample forecasts. Not surprising, the in-sample forecasts from the stationary MS-AR(1) model are always superior to forecasts from both the single regime RW model and the MS-RW model.

Table 3: Out-of-sample mean squared forecast errors of a random walk model and improvements (percent) of a Markov switching random walk model (MS-RW) and a Markov switching AR(1) model (MS-AR(1)).

Country	Model	$\ln(L)$	Forecast horizon			
			1	2	3	4
UK	Random walk	-220.5	25.3	59.8	77.0	97.3
	MS-RW	-217.1	-7.7%	-14.8%	-15.8%	-17.0%
	MS-AR(1)	-214.9	6.0%	9.1%	16.5%	23.0%
France	Random walk	-215.3	21.2	51.3	73.1	100.7
	MS-RW	-209.0	5.2%	0.1%	0.5%	-5.9%
	MS-AR(1)	-202.7	7.2%	13.9%	19.3%	20.4%
Germany	Random walk	-220.3	22.6	53.7	74.7	105.9
	MS-RW	-216.7	4.9%	0.2%	3.1%	-1.1%
	MS-AR(1)	-212.0	7.8%	12.5%	19.1%	22.1%
Switzerland	Random walk	-225.8	29.6	68.1	97.9	147.3
	MS-RW	-218.3	-3.3%	-7.1%	4.3%	-1.0%
	MS-AR(1)	-217.7	4.1%	7.3%	24.7%	22.8%
Canada	Random walk	-133.5	4.0	10.1	18.8	31.8
	MS-RW	-127.8	13.6%	17.7%	19.2%	14.9%
	MS-AR(1)	-125.6	11.6%	15.9%	17.7%	15.8%
Japan	Random walk	-221.5	27.2	62.4	93.5	138.8
	MS-RW	-217.9	-4.3%	-8.1%	-2.2%	1.7%
	MS-AR(1)	-213.9	11.4%	7.3%	3.6%	1.7%

Note: The population parameters were estimated using the effective sample 1973:2 to 1990:4 and mean squared errors are those associated with forecast errors for dates 1991:1+m to 1997:4, where m is the forecast horizon.

passes the RW model for all forecast horizons, but we also reject the null the RW model encompasses the MS-AR(1) model for two forecast horizons. Based on these results, we conclude that forecasts using the MS-AR(1) model are significantly better than forecasts using the two competing models for five of the six real exchange rates in our sample. The results for the Japanese Yen are inconclusive as the tests both suggest that the MS-AR(1) model encompasses the RW model and vice versa.

The forecasting behavior of our model can also be compared to findings in Siddique and Sweeney (1998) who compares out-of-sample forecasts from a stationary AR(1) model with random walk forecasts. They show that forecasts of real exchange rates were dependent on the size of the estimation sample, such that longer samples tend to produce better forecasts than when using shorter estimation periods. When using long samples, out-of-sample forecasts may out-perform forecasts from a random walk model. By contrast, for short estimation periods, it may be the case that one state dominates the sample leading to a biased estimate of the parameters of the model and hence, forecasts tend to be worse. Siddique and Sweeney (1998) also found that improvements of forecasts were positively related to the forecast horizon. The results in Table 3 also confirm this finding for forecasts generated by the stationary Markov switching AR(1) model.

3.4 Rejection frequencies of ADF-tests

An alternative approach to evaluate our base line model is to examine whether the Markov model can explain the failure of not rejecting unit roots in post-Bretton Woods data. For this experiment, we assume that real exchange rates are generated by our base model, but that the econometrician assumes that there are no structural breaks in the parameters. How often will this econometrician reject a unit root in the real exchange rate when using standard unit root tests such as the ADF-test?

To answer this question, we use the estimated parameters shown in Table 1 (using the sample period 1973:2-1990:4) to generate 10000 samples of artificial observations of the six real exchange rates. For each sample, we first construct an artificial sample of the state process using the estimate of the Markov transition probabilities, \hat{P} , where the first observation of the state process is obtained using the ergodic probabilities. Then, we construct a sequence of residuals by drawing normally distributed random numbers with variances given in Table 1, σ^2 . Using the actual estimated parameters, μ_i and ϕ , we then generate each sample of artificial real exchange rates. For these 10000 samples, we finally calculate the ADF-test. Since our base model is stationary, we expect to reject unit roots quite frequently given that the ADF-test has power against nonlinearity of this form.

Table 5, reports rejection frequencies at the 5 and 10 percent levels of ADF-tests using 3 lags, when the true data generating process is a stationary Markov switching AR(1) model. In general, the probability of rejecting unit roots in the simulated data is very low at conventional significance levels. In more than 85 percent of the cases, we

Table 4: Modified Diebold–Mariano tests of out-of-sample forecast encompassing of two Markov switching models (MS-RW and MS-AR(1)) over a single regime random walk model (RW). Only p-values are reported.

Country	Nullhypothesis	Forecast horizon			
		1	2	3	4
UK	RW \mathcal{E} MS-RW	0.645	0.920	0.998	0.987
	MS-RW \mathcal{E} RW	0.114	0.019	0.000	0.003
	RW \mathcal{E} MS-AR(1)	0.099	0.109	0.087	0.061
	MS-AR(1) \mathcal{E} RW	0.535	0.600	0.661	0.711
France	RW \mathcal{E} MS-RW	0.065	0.167	0.162	0.485
	MS-RW \mathcal{E} RW	0.305	0.155	0.211	0.060
	RW \mathcal{E} MS-AR(1)	0.030	0.014	0.017	0.020
	MS-AR(1) \mathcal{E} RW	0.266	0.558	0.719	0.700
Germany	RW \mathcal{E} MS-RW	0.065	0.265	0.189	0.360
	MS-RW \mathcal{E} RW	0.432	0.273	0.437	0.244
	RW \mathcal{E} MS-AR(1)	0.011	0.029	0.043	0.062
	MS-AR(1) \mathcal{E} RW	0.339	0.485	0.536	0.534
Switzerland	RW \mathcal{E} MS-RW	0.297	0.447	0.180	0.326
	MS-RW \mathcal{E} RW	0.148	0.106	0.391	0.271
	RW \mathcal{E} MS-AR(1)	0.130	0.112	0.066	0.113
	MS-AR(1) \mathcal{E} RW	0.281	0.331	0.695	0.641
Canada	RW \mathcal{E} MS-RW	0.025	0.065	0.090	0.127
	MS-RW \mathcal{E} RW	0.696	0.686	0.705	0.651
	RW \mathcal{E} MS-AR(1)	0.047	0.084	0.098	0.103
	MS-AR(1) \mathcal{E} RW	0.739	0.705	0.685	0.643
Japan	RW \mathcal{E} MS-RW	0.476	0.608	0.396	0.268
	MS-RW \mathcal{E} RW	0.194	0.099	0.248	0.390
	RW \mathcal{E} MS-AR(1)	0.029	0.082	0.102	0.120
	MS-AR(1) \mathcal{E} RW	0.045	0.008	0.000	0.000

Note: The population parameters were estimated using the effective sample 1973:2 to 1990:4 and mean squared errors are those associated with forecast errors for dates 1991:1+m to 1997:4, where m is the forecast horizon. X \mathcal{E} Y denotes that forecasts from model X encompasses forecasts from model Y.

Table 5: Rejection probabilities (in percent) of ADF-tests at different significance levels when the true data generating process is a 2-state Markov switching AR(1) model, 10000 trials.

	UK	France	Germany	Switzerland	Canada	Japan
	Rejection probability					
5% level	9.8%	11.5%	13.3%	6.9%	5.4%	12.0%
10% level	18.5%	22.2%	24.4%	14.0%	10.8%	23.2%
Mean ADF	-2.00	-2.10	-2.17	-1.81	-1.64	-2.13
Sample ADF	-1.78	-1.72	-1.86	-2.08	-2.05	-1.68

Note: Mean ADF is the average value of the ADF t-ratio for the simulated data. Sample ADF refers to the ADF t-ratio for the actual data shown in Figure 1. All ADF-tests are calculated using 3 lags.

cannot reject unit roots at the 5 percent level. However, the results also suggest that when persistence is relatively low, the probability of rejecting the null is relatively high. For example, the degree of persistence (measured as the maximum eigenvalue of the matrix A , $\max |\text{eig}(\hat{A})|$ in Table 1) for Japanese Yen 0.759 and we reject the null in 12 percent of the cases. By contrast, for the Swiss Franc where the autoregressive parameter is relatively close to unity (0.918), we reject a unit root in about 7 percent of the cases using the five percent significance level.

These results are also consistent with tests of unit roots in the actual data reported in the fourth row of Table 5. We cannot reject unit roots in any of the six real exchange rates in our sample at conventional significance levels. Our reading of these results is that the ADF test cannot distinguish between a stationary 2-state Markov model and a unit root model. This is similar to results reported in Perron and Vogelsang (1992) where the distribution of the ADF test is affected by a shift in the intercept. The failure of rejecting unit roots when assuming no state dependence is therefore not surprising if the true model is a stationary Markov switching model. If the econometrician assumes that there are no regime switches, and applies the ADF-test, the evidence suggests that real exchange rates are non-stationary even when the true data generating process is a stationary MS-AR(1) model.

4 Conclusions

This paper suggests that the real exchange rate can be modelled as a stationary 2-state Markov switching AR(1) model. Following the tradition in the literature (cf. Meese

and Rogoff (1983)) we use the out-of-sample forecasting behavior as our main indicator to discriminate between different empirical models. Using the mean squared error to distinguish between out-of-sample forecasts, we find that our base model, where only the intercept varies across regimes, out-predicts both the single regime random walk and the Markov switching random walk models. Formal tests of forecast encompassing also show that the forecasts using our base model are significantly better than forecasts using either a one state or a two-state random walk model for all currencies except for the Japanese Yen. This is surprising since exchange rate models often are outperformed by random walk models.

A common finding when testing whether post-Bretton Woods real exchange rates contain unit roots is that this null cannot be rejected. Given these empirical findings and if our base model describes the data generating process of real exchange rates, we expect that unit root tests such as the ADF-test should have low power against this particular form of nonlinear serial dependence. Our simulation experiments support this view. Similarly to earlier findings in Perron (1990) who examined the power of unit root tests when there is a single break in the intercept, we find that the ADF-test is not capable of distinguishing between unit roots and multiple shifts in the intercept of our Markov switching model. Therefore, it is not surprising that ADF-tests on actual post-Bretton Woods data usually fail to reject the null of unit roots in real exchange rate data. Our main conclusion from this empirical analysis (simulations as well as out-of-sample forecasts) is that the stationary Markov switching AR(1) model allowing only the intercept to vary over time out-performs two competing random walk models.

References

- Cheung, Y.-W. and Lai, K. S. (1993a). A Fractional Cointegration Analysis of Purchasing Power Parity. *Journal of Business and Economic Statistics*, 11:103–12.
- Cheung, Y.-W. and Lai, K. S. (1993b). Long-Run Purchasing Power Parity during the Recent Float. *Journal of International Economics*, 34:181–92.
- Creedy, J., Lye, J., and Martin, V. L. (1996). A Non-linear Model of the Real US-UK Exchange Rate. *Journal of Applied Econometrics*, 11:669–686.
- Dropsy, V. (1996). Real Exchange Rates and Structural Breaks. *Applied Economics*, 28:209–219.
- Dumas, B. (1992). Dynamic Equilibrium and the Real Exchange Rate in a Spatially Separated World. *Review of Financial Studies*, 5:153–180.

- Edison, H. J., Gagnon, J. E., and Melick, W. R. (1997). Understanding the Empirical Literature on Purchasing Power Parity: The Post-Bretton Woods Era. *Journal of International Money and Finance*, 16:1–17.
- Engel, C. (1994). Can the Markov Switching Model Forecast Exchange Rates? *Journal of International Economics*, 36:151–165.
- Engel, C. (2000). Long-Run PPP May Not Hold After All. *Journal of International Economics*, forthcoming.
- Engel, C. and Hamilton, J. D. (1990). Long Swings in the Dollar: Are They in the Data and Do Markets Know It? *American Economic Review*, 80:689–713.
- Engel, C. and Kim, C.-J. (1999). The Long-Run U.S./U.K. Real Exchange Rate. *Journal of Money, Credit, and Banking*, 31:335–356.
- Fritsche, C.-P. and Wallace, M. (1997). Forecasting the Exchange Rate PPP versus a Random Walk. *Economics Letters*, 54:69–74.
- Froot, K. A. and Rogoff, K. (1995). Perspectives on PPP and Long-Run Real Exchange Rates. In Grossman, G. and Rogoff, K., editors, *Handbook of International Economics, Vol. III*. North-Holland, Amsterdam.
- Hamilton, J. D. (1996). Specification Testing in Markov-Switching Time-Series Models. *Journal of Econometrics*, 70:127–157.
- Hansen, B. E. (1992). The Likelihood Ratio Test under Non-Standard Conditions: Testing the Markov Switching Model of GNP. *Journal of Applied Econometrics*, 7:S61–S82.
- Hansen, B. E. (1996). Erratum: The Likelihood Ratio Test Under Nonstandard Conditions: Testing the Markov Switching Model of GNP. *Journal of Applied Econometrics*, 11:195–198.
- Harvey, D., Leybourne, S., and Newbold, P. (1998). Test for Forecast Encompassing. *Journal of Business and Economic Statistics*, 16:254–259.
- Holst, U., Lindgren, G., Holst, J., and Thuvsholmen, M. (1994). Recursive Estimation in Switching Autoregressions with a Markov Regime. *Journal of Time Series Analysis*, 15:489–506.
- Jorion, P. and Sweeney, R. J. (1996). Mean Reversion in Real Exchange Rates: Evidence and Implications for Forecasting. *Journal of International Money and Finance*, 15:535–550.

- Juselius, K. (1995). Do Purchasing Power Parity and Uncovered Interest Rate Parity Hold in the Long Run? An Example of Likelihood Inference in a Multivariate Time-Series Model. *Journal of Econometrics*, 69:211–40.
- Karlsen, H. (1990). *A Class of Non-Linear Time Series Models*. PhD thesis, Department of Mathematics, University of Bergen, Norway.
- Lothian, J. R. (1997). Multi-Country Evidence on the Behavior of Purchasing Power Parity under the Current Float. *Journal of International Money and Finance*, 16:19–35.
- Meese, R. A. and Rogoff, K. (1983). Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample? *Journal of International Economics*, 14:3–24.
- Norrbin, S. C., Reffett, K. L., and Ji, Y. (1997). Using a VECM to Test Exogeneity and Forecastability in the PPP Condition. *Applied Financial Economics*, 7:87–95.
- Oh, K.-Y. (1996). Purchasing Power Parity and Unit Root Tests Using Panel Data. *Journal of International Money and Finance*, 15:405–18.
- Ott, M. (1996). Post Bretton Woods Deviations from Purchasing Power Parity in G7 Exchange Rates — An Empirical Exploration. *Journal of International Money and Finance*, 15:899–924.
- Perron, P. (1990). Testing for a Unit Root in a Time Series With a Changing Mean. *Journal of Business and Economic Statistics*, 8:153–162.
- Perron, P. and Vogelsang, T. J. (1992). Nonstationarity and Level Shifts with an Application to Purchasing Power Parity. *Journal of Business and Economic Statistics*, 10:301–20.
- Rogoff, K. (1996). The Purchasing Power Parity Puzzle. *Journal of Economic Literature*, 34:647–68.
- Sarantis, N. (1999). Modelling Non-Linearities in Real Effectiv Exchange Rates. *Journal of International Money and Finance*, 18:27–45.
- Siddique, A. and Sweeney, R. J. (1998). Forecasting Real Exchange Rates. *Journal of International Money and Finance*, 17:63–70.