
Can Parameter Instability Explain the Meese-Rogoff Puzzle?

Philippe Bacchetta, *University of Lausanne and CEPR*

Eric van Wincoop, *University of Virginia and NBER*

Toni Beutler, *University of Lausanne and Study Center Gerzensee*

I. Introduction

The empirical literature on nominal exchange rates shows that the current exchange rate is often a better predictor of future exchange rates than a linear combination of macroeconomic fundamentals. This result is behind the famous Meese-Rogoff puzzle. In their seminal work, Meese and Rogoff (1983a, 1983b) estimate linear regression models based on standard macroeconomic variables. Using rolling regressions, they show that forecasts based on these models do not outperform forecasts based on the current exchange rate, even when the actual *future* macro fundamentals are used. Their results have largely held up since then, even with much more data available (see, e.g., Cheung, Chinn, and Pascual 2005; Rogoff and Stavrakeva 2008). A potential explanation of this puzzle is that the relationship between nominal exchange rates and macroeconomic fundamentals is unstable. There is widespread evidence documenting this instability.¹ In their original work, Meese and Rogoff themselves already conjectured that parameter instability may explain their results.

The goal of this paper is to evaluate whether parameter instability can indeed account for the Meese-Rogoff puzzle. In order to do so, we proceed in three steps. We first conduct a Meese-Rogoff exercise on five currencies of industrialized countries from 1975 to 2008. We estimate rolling regressions and forecast out of sample using actual future fundamentals. We then compute the mean square prediction error (MSE) and compare it with the MSE resulting from a prediction based on the current exchange rate. Our results confirm once again the original Meese-Rogoff findings: exchange rate depreciations are better predicted by a random walk than by the estimated linear model. In the second step, we assume a theoretical reduced-form relationship between exchange rates and

fundamentals in which parameters are constant. We calibrate the model to data for exchange rates and fundamentals for the five currencies. In the final step, we introduce exogenous parameter instability to the relationship between exchange rates and fundamentals.² We then conduct the exact same Meese-Rogoff exercise with the data generated by the reduced-form model, for both constant and time-varying parameters. To help interpret the findings, we also derive theoretical results on the impact of time-varying parameters on the out-of-sample forecasting performance of the model.

It is easy to see why it is natural to consider time-varying parameters in accounting for the Meese-Rogoff puzzle. If parameters were constant and known, the linear model would by construction outperform the random walk. As long as the observed macro fundamentals have any explanatory power, the model obviously has more explanatory power than a random walk forecast. In order to explain the Meese-Rogoff findings, we therefore have to relax the assumption that parameters are constant and known. One way to do this is to assume that parameters are constant but not known. With samples of finite length, parameters are estimated with error. Such estimation error contributes to a forecasting error and can explain the Meese-Rogoff findings. Not surprisingly, this has received significant attention in the literature, and statistics have been developed to correct for such small-sample bias (e.g., Clark and West 2006).

But the forecasting performance can further deteriorate when parameters themselves are varying over time. Even if we ignore the small-sample estimation errors, a finite sample provides an estimation of a weighted average of parameters over the estimation sample. This average of past parameters is not necessarily a good measure of future parameters. The resulting further deterioration is what Meese and Rogoff had in mind when pointing to time-varying parameters as a possible resolution to the weak out-of-sample performance of the model. However, we find that time-varying parameters also work in another, opposite, direction. Abstracting for a moment from estimation errors of the parameters, we show that time variation in parameters improves the average explanatory power of fundamentals. The reason is that parameters sometimes become high in absolute value, and therefore, fundamentals have more explanatory power. This second implication of time-varying parameters actually improves the out-of-sample performance of the model relative to the random walk.

We find that the two effects typically offset each other when the reduced-form model is calibrated to the data. Thus, the impact of

time-varying parameters on prediction performance is very small. We show that there are two cases in which the impact of time-varying parameters on the out-of-sample performance can become significant, but neither can explain the Meese-Rogoff puzzle. One is the case in which the persistence of parameters is close to one. But in this case time variation implies a *better* prediction performance, so this goes in the wrong direction. The second case is one in which fundamentals have high explanatory power. In this case parameter instability can substantially deteriorate the out-of-sample performance of the model. But in reality the observed fundamentals have very limited explanatory power. More important, in this case there is no Meese-Rogoff puzzle because the model always outperforms the random walk, whether parameters vary or not. We conclude that it is not time-varying parameters but rather small-sample estimation bias that explains the Meese-Rogoff puzzle.

The remainder of the paper is organized as follows. In Section II we discuss empirical results from the Meese-Rogoff exercise. In Section III we propose a theoretical reduced-form exchange rate model, with either constant or time-varying parameters. We calibrate the model with constant parameters to the data. In Section IV we discuss results from conducting the Meese-Rogoff exercise on the model. We find that time-varying coefficients play almost no role. To shed further light on these findings, in Section V we derive theoretical results on the impact of time-varying parameters on the ability of the model to forecast out of sample. We then connect these results to the findings from the simulations in Section IV. In Section VI, we extend this analysis to the in-sample fit. Section VII examines some further implications of the theoretical model. Section VIII presents conclusions.

II. Out-of-Sample and In-Sample Fit in the Data

In order to evaluate both the out-of-sample and in-sample relationships between exchange rates and fundamentals, we consider five currencies relative to the U.S. dollar: Swiss franc, British pound, Canadian dollar, Japanese yen, and German mark (euro since 1999). We use monthly data from September 1975 to September 2008. The five macro fundamentals that we consider as exchange rate predictors are standard: differential of money supply growth, industrial production growth, and unemployment rate growth relative to the United States; growth in the oil price; and the lagged interest rate differential relative to the United States.³ Following Meese-Rogoff and most of the literature, the regressions

are estimated individually.⁴ A precise description of the data and data sources can be found in appendix A.

Figure 1 reports the relative out-of-sample fit for each of the five currencies, as well as the average across the five currencies. It is the ratio of the MSE of a one-period-ahead forecast from the estimated model relative to the MSE of a random walk (or no change) forecast. The model forecasts are based on rolling regressions of sample length L . The first regression is run on a sample of length L that starts in September 1975. After regressing the change in the log exchange rate on the five macro fundamentals over this sample (plus a constant), we forecast 1 month out of sample using the estimated parameters of the fundamentals together with the *actual* macro fundamentals 1 month out of sample.⁵ The difference between the “forecast” of the change in the log exchange rate 1 month out of sample and the actual change in the log exchange rate is the forecast error. This is subsequently repeated for a sample of length L that starts 1 month later, in October 1975, and so on. We conduct a total of $P = 200$ rolling regressions in order to compute the MSE. For the random walk the forecast error is the actual change in the log exchange rate since the forecasted change is zero. The standard measure of relative out-of-sample fit of the model is the ratio of the MSE of the model to the MSE of the random walk.⁶

Several conclusions can be drawn from figure 1. First, it shows that the Meese-Rogoff finding that the model does not outperform the random walk continues to hold up in the data. With the exception of $L > 150$ for Canada, the model fails to outperform the random walk for all currencies. Second, as expected, the relative performance of the model improves with the sample length as estimation error of the parameters becomes less severe for longer samples. The average for the five currencies shows that the MSE ratio gradually drops from 1.21 for $L = 40$ to 1.02 for $L = 196$ (the maximum sample length). But it is remarkable that even for the relatively long sample of $L = 196$, which is 16.3 years, the model on average still does not outperform the random walk.

A final point to notice about figure 1 is that there are significant differences across currencies. The MSE ratio decreases gradually with a rise in L for the Swiss franc, German mark, and Canadian dollar, but it does not show a strong trend for the yen and it suddenly rises around $L = 100$ for the British pound. Also, in contrast to the other currencies, the Canadian dollar is the only one for which the model does outperform the random walk for a long enough horizon. Such differences are to be expected as estimation error of the parameters depends on the specific shocks that hit these currencies during the sample. As emphasized

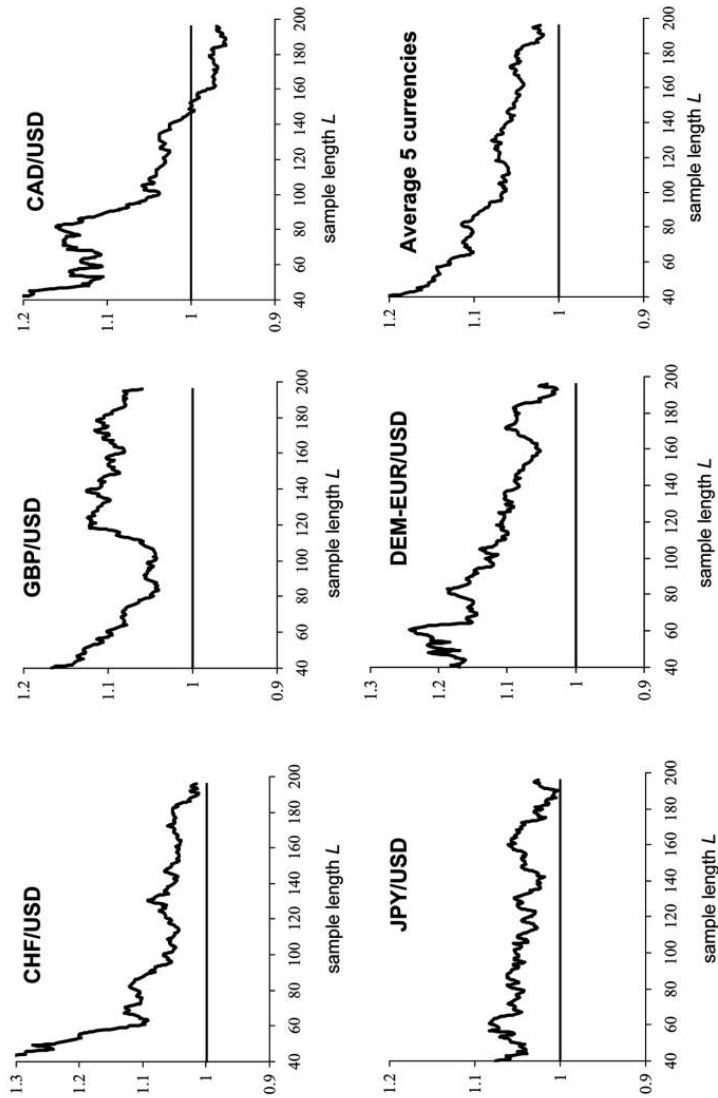


Fig. 1. Out-of-sample fit in data: $\text{MSE}^{\text{MODEL}} / \text{MSE}^{\text{RW}}$. Each chart reports the ratio of the MSE of 1-month-ahead exchange rate forecasts from the model relative to the MSE of the random walk. Forecasts from the model are based on rolling regressions of sample length L (horizontal axis). The model includes the following regressors: differential with the United States of money supply growth, industrial production growth, and unemployment rate growth; growth in the oil price; and the lagged interest rate differential in levels. The MSEs are computed on $P = 200$ forecasts. The results are reported for five currencies. The last chart reports the average MSE ratio over the five currencies. The sample is 1975:9–2008:9.

by Cheung et al. (2005) or Alquist and Chinn (2008), no model consistently outperforms the random walk by the MSE criterion. Some models (i.e., some sets of explanatory fundamentals) can outperform the random walk for some currency (as in our case for Canada when the sample is long enough), but no model consistently outperforms the random walk across currencies and samples.

Figure 2 reports both the out-of-sample fit (same as fig. 1) and the in-sample fit. The latter is one minus the average in-sample R^2 . For a particular sample length L , the average R^2 is computed as the average R^2 over the same $P = 200$ rolling regressions that are used to estimate parameters for the out-of-sample forecasts. Figure 2 confirms what is well known, that the in-sample fit is better than the out-of-sample fit. There would be no difference between the two if parameters were known (estimated without error), whether they are constant or time varying. It is the estimation error of parameters that causes the in-sample fit to be better than the out-of-sample fit. Estimation error reflects a spurious fit within sample, causing the R^2 to be particularly high for low L . But at the same time it is a source of forecast errors in the out-of-sample exercise that deteriorates the performance of the model relative to the random walk by the MSE criterion. This is illustrated nicely in the last chart, for the average of the five currencies. While the out-of-sample fit deteriorates as L decreases (MSE ratio rises), the in-sample fit improves ($1 - R^2$ goes down).

The average R^2 across all five currencies ranges from 0.16 for $L = 40$ to 0.04 for $L = 196$. There are significant differences across currencies. For $L = 196$ the R^2 ranges from 0.02 for the German mark to 0.06 for the yen. There is no straightforward relationship between in- and out-of-sample performance. One might expect that currencies for which fundamentals have more explanatory power have both a better in- and out-of-sample performance. But we have already seen that small-sample estimation error improves the in-sample fit whereas it deteriorates the out-of-sample fit. This may explain why, for example, the out-of-sample fit for $L = 196$ is much better for the Canadian dollar than for the yen (MSE ratio of 0.97 vs. 1.03), whereas the in-sample fit is worse for the Canadian dollar than for the yen (R^2 of 0.05 vs. 0.06).

In the context of the rolling regressions, it is interesting to see that parameter estimates move over time. This could potentially be an indication of time variation. This time variation is illustrated in figure 3, for a specific coefficient. In the three charts of that figure, we report the estimated coefficients associated with the money growth differential in the Japanese yen/U.S. dollar rolling regressions. The three charts correspond

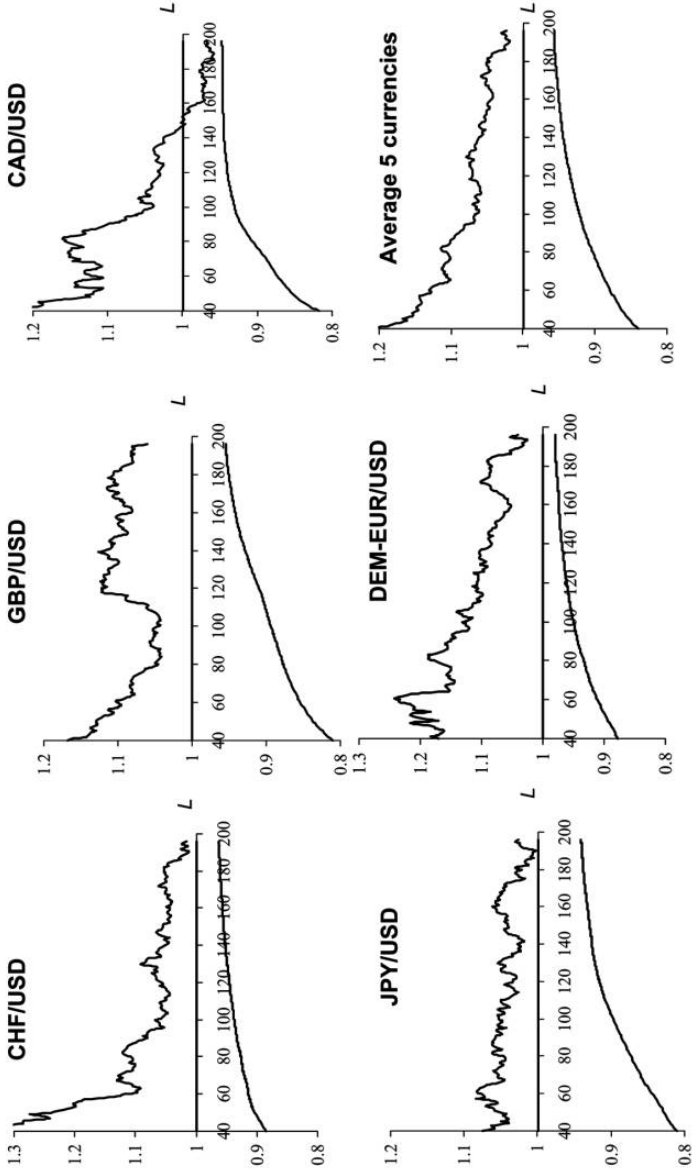


Fig. 2. In-sample versus out-of-sample fit in data: $1 - R^2$ versus $\text{MSE}^{\text{MODEL}}/\text{MSE}^{\text{RW}}$. Each chart reports the ratio of the MSE of 1-month-ahead exchange rate forecasts from the model relative to the MSE of the random walk (top line) and $1 - R^2$ (bottom line). Forecasts from the model are based on rolling regressions of sample length L (horizontal axis). The model includes the following regressors: differential with the United States of money supply growth, industrial production growth, and unemployment rate growth; growth in the oil price; and the lagged interest rate differential in levels. The MSEs are computed on $P = 200$ forecasts. The R^2 is computed as the average in-sample R^2 across the rolling regressions. The results are reported for five currencies. The last chart reports the average over the five currencies. The sample is 1975:9–2008:9.




Fig. 3. Estimated coefficients from rolling regressions in data (JPY/USD: relative money supply growth). The three charts report estimated coefficients from rolling regressions over three different regression sample lengths L . The variable considered is the differential in money supply growth in the Japanese yen/U.S. dollar exchange rate model. The x axis indicates the date of the first observation in the regression sample.

to three different sample lengths: $L = 40, 120,$ and 200 .⁷ The first observation on each chart is the value of the estimated regression coefficient over a sample that starts in September 1975 and contains L data points. We then shift the whole estimation sample one period and estimate the coefficient again, and we repeat the procedure until we reach the end of the sample. As we would expect, the estimated coefficients appear more time varying for smaller regression samples. For $L = 40$, the coefficient varies from -8 to $+10$, whereas it varies only from -1.5 to $+4$ for $L = 200$.

Finally, it is interesting to note that even though most exchange rate models cannot beat the random walk, there are some exceptions. One is the case of commodity currencies.⁸ Figure 4 shows the average MSE




Fig. 4. Out-of-sample fit for commodity currencies. This chart reports the average ratio of the MSE of 1-month-ahead exchange rate forecasts from the model relative to the MSE of the random walk over three commodity currencies: Australian, Canadian, and New Zealand dollar versus the U.S. dollar. Forecasts from the model are based on rolling regressions of sample length L (horizontal axis). The model includes the contemporaneous change in the log of the country-specific index of commodity prices and a constant as regressors. The MSEs are computed on $P = 120$ forecasts. The results are based on the sample 1986:1–2008:12.

ratio across three currencies against the U.S.dollar: Australian, Canadian, and New Zealand dollars. For each currency, the MSE ratio is computed using the same procedure as above, but over $P = 120$ forecasts because of the shorter sample available.⁹ One-month-ahead forecasts from the model are based on a regression of the change in the log nominal exchange rate on the contemporaneous change in the log of the country-specific index of commodity prices (and a constant). The average MSE ratio is below one for every sample size, so that the model clearly beats the random walk.

III. Model and Calibration

We will adopt the following reduced-form exchange rate model:

$$\Delta s_t = \sum_{n=1}^N \beta_{nt} f_{nt} + u_t, \quad (1)$$

$$f_{nt} = \rho_n f_{n,t-1} + \epsilon_{nt}^f, \quad (2)$$

$$u_t = \rho_u u_{t-1} + \epsilon_t^u, \quad (3)$$

$$\beta_{nt} - \beta_n = \rho_\beta (\beta_{n,t-1} - \beta_n) + \epsilon_{nt}^\beta, \quad (4)$$

where s_t is the log exchange rate and f_{nt} represents fundamental n . The constant parameter case corresponds to $\rho_\beta = 0$ and $\epsilon_{nt}^\beta = 0$ for all n, t . The innovations are normally distributed with mean zero and variance

$$\text{Var} \begin{pmatrix} \epsilon_t^f \\ \epsilon_t^u \\ \epsilon_t^\beta \end{pmatrix} = \begin{pmatrix} \Omega_f & 0 & 0 \\ 0 & \sigma_u^2 & 0 \\ 0 & 0 & \sigma_\beta^2 I_N \end{pmatrix}, \quad (5)$$

where

$$\epsilon_t^f = \begin{pmatrix} \epsilon_{1t}^f \\ \dots \\ \epsilon_{Nt}^f \end{pmatrix}, \quad \epsilon_t^\beta = \begin{pmatrix} \epsilon_{1t}^\beta \\ \dots \\ \epsilon_{Nt}^\beta \end{pmatrix}. \quad (6)$$

The change in the exchange rate depends on N observed fundamentals f_{nt} . It is also driven by unobserved fundamentals, which are summarized in u_t . The fundamentals follow AR(1) processes, with generally different autoregressive coefficients. We allow for a general variance-covariance structure of the innovations in the observed fundamentals.

The unobserved fundamental u_t also follows an AR(1) process. Its innovation is uncorrelated with that of the observed fundamentals. When allowing for parameter uncertainty, we assume an AR(1) process for each of the parameters. We assume that the parameter innovations are uncorrelated across fundamentals.

In calibrating the model, we match the key moments of fundamentals and exchange rates in the data. This is done as follows. For each of the five currencies in Section II, we regress Δs_t on a constant and each of the five observed fundamentals, using the entire sample of monthly data from September 1975 to September 2008. For fundamental n this gives us an estimate of the mean parameter value β_n . Without loss of generality we then redefine the fundamentals, by multiplying them with appropriate constants, in order to normalize all β_n to one. For example, when the estimated coefficient is 0.5, we define a new fundamental that is 0.5 times the old fundamental. The estimated coefficient for the new fundamental is then one, which is our estimate for β_n . This procedure has the advantage that all fundamentals have the same mean coefficients.

Next we estimate the AR(1) processes (2) for the fundamentals for each of the currencies. For fundamental n we set the AR coefficients ρ_n equal to the average of the estimated AR(1) coefficients across the five currencies. We use the estimated innovations to compute the correlation matrix of fundamental innovations as well as their standard deviations. We then set the correlation matrix for the innovations equal to the average across the five currencies and similarly for the standard deviation of the fundamental innovations. These numbers give us the matrix Ω_f .

One comment is in order about this procedure so far. When applying the model to each of the five currencies of Section II, we will assume the same AR coefficients ρ_n and covariance matrix Ω_f for each of the currencies. One can also use separate estimates for each currency. The disadvantage of that approach, though, is that the results will very much depend on estimates of the mean parameters for each of the currencies (before they are normalized to one), which are subject to small-sample estimation error that will affect the standard deviation of the normalized fundamentals. Similarly, small-sample errors will affect the estimate of ρ_n for individual currencies. To minimize such errors, we average across currencies to compute ρ_n and the standard deviation of fundamental innovations. We find that the results from model simulation using this procedure fit the data better than estimating the ρ_n and Ω_f separately for each currency.

Finally, we need an estimate of the standard deviation and persistence ρ_u of the error u_t in the exchange rate equation. We estimate this separately for each currency by matching the observed standard deviation and first-order autocorrelation of Δs_t . We do so for the constant parameter case, but the results are virtually identical for the time-varying parameter case since overall exchange rate volatility is not much affected by time-varying parameters.¹⁰

We will not use data to estimate σ_β and ρ_β for the time-varying parameter case. This is related to the key finding of the paper: it is very hard to empirically distinguish between constant and time-varying parameters. Therefore, instead we consider a wide range of assumptions about ρ_β , and in most of the analysis we will set σ_β such that the unconditional standard deviation of the parameters is quite large: equal to the mean value one of the parameters.¹¹

IV. Impact of Time-Varying Parameters

In this section we use the model presented above to generate data and compute the MSE ratios as in Section II. Before considering the impact of time-varying parameters on both the out-of-sample and in-sample fit, we first discuss the results from model simulations for the constant parameter case. For each currency we conduct 1,000 simulations of the model over 397-month samples, corresponding to the September 1975 to September 2008 sample in the data. For each simulation, we first generate a history of 1,000 months prior to our 397-month sample in order to avoid having to start from a steady state. All innovations are drawn from the normal distributions discussed in the previous section.

Figure 5 reports for each currency the MSE ratio in the model relative to the random walk. Results are reported for both the data and the model. For the model, there are three lines: the average over the 1,000 simulations and the upper and lower bands of the 99% confidence interval based on the 1,000 simulations. The results for the data generally conform to those for the model simulations. The MSE ratio in the data generally falls within the 99% confidence band for the model. There are a few exceptions when it rises slightly above the confidence band. Particularly noteworthy are the German mark, for which the data are generally close to the upper band of the confidence interval, and the British pound, for which the data are pretty much on top of the upper band of the confidence interval for $L > 120$. As is the case in the data, the average MSE ratio in the model (across the 1,000 simulations) remains

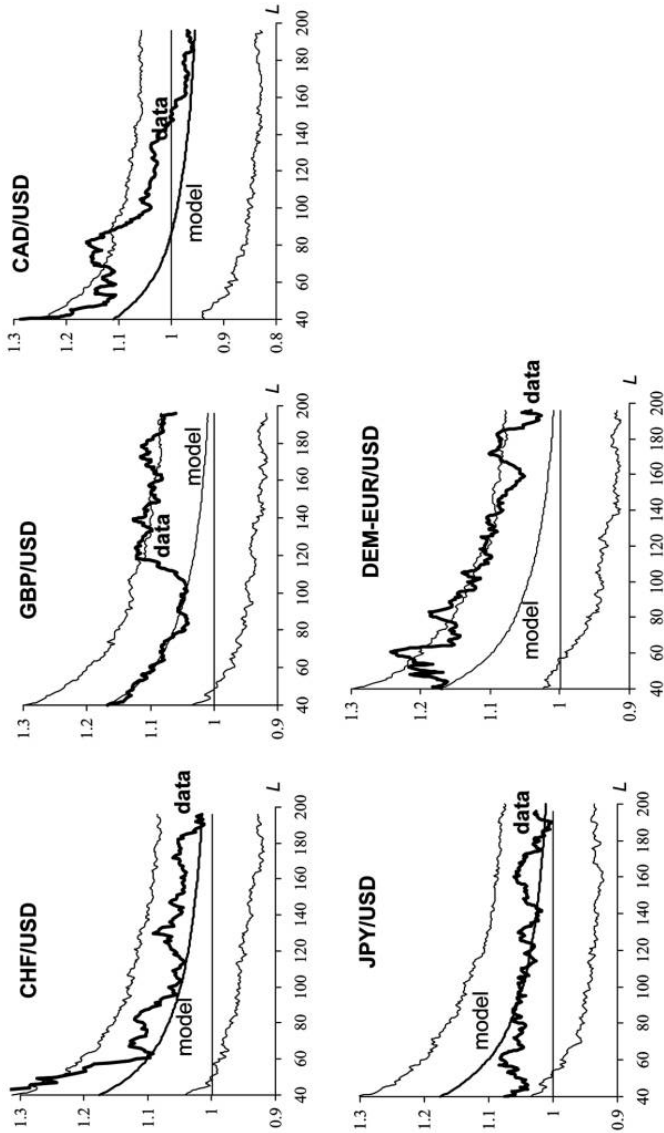


Fig. 5. Out-of-sample fit model versus data: MSE^{MODEL}/MSE^{RW} . Each chart contains four lines. The bold line represents the out-of-sample fit in the data and is the same as in figures 1 and 2: MSE of 1-month-ahead exchange rate forecasts including money, output, unemployment rate, oil price, and lagged interest rate estimated by rolling regressions relative to the MSE of random walk forecast. The horizontal axis shows the sample length L of the rolling regressions. The results are based on $P = 200$ rolling regressions. The downward-sloping smooth thin line is the corresponding statistic in the constant parameter model, computed as an average over 1,000 simulations of the model for the same L and P as in the data. Finally, the top and bottom thin lines represent 99% confidence intervals based on the 1,000 simulations of the model.

above one for all currencies but the Canadian dollar, for which it reaches below one for L sufficiently big. Also, as is the case for the average of the currencies, the average MSE ratio in the model gradually falls as L rises.

Figure 6 reports the results for the in-sample fit. There are again four lines, which represent the data, the average of 1,000 simulations of the model, and the upper and lower bands of the 99% confidence interval based on the 1,000 simulations. The results for the data always lie within the confidence band based on the model simulations. In both the model simulations and the data, the average R^2 always declines in the sample length L . For the Swiss franc, Canadian dollar, and German mark, the R^2 is somewhat lower in the data than the average over 1,000 model simulations. For the other two currencies, the British pound and the yen, it is the other way around.

Overall the data are consistent with the model simulations for the constant parameter case. This may lead one to believe that time-varying parameters are not needed to explain the in- and out-of-sample fit. To examine this more closely, we will now compare the in- and out-of-sample fit for the constant parameter case to the time-varying parameter case. This is done in figures 7 and 8 for $\rho_\beta = 0$, which report, respectively, the average out-of-sample and average in-sample fit across 1,000 simulations of both the constant and time-varying parameter cases. Different values for ρ_β are considered in figures 9 and 10. In each of these cases the unconditional standard deviation of the parameters is set at one, equal to the mean value of the parameters. This amounts to considerable variation over time in the parameters. A two-standard-deviation band of the parameters ranges from -1 to $+3$.

Figure 7 reports results for the MSE ratio for each of the currencies, as well as the average across the currencies. Each chart contains two lines, which represent the MSE ratio for the constant (solid line) and time-varying parameter (dotted line) cases. With the marginal exception of the Canadian dollar, the MSE ratio for the constant parameter case is virtually indistinguishable from the ratio for the time-varying parameter case. The two lines are virtually on top of each other. In order to show that there is a slight difference, in the bottom charts of figure 7 we zoom in for L from 120 to 130, with a much narrower range of numbers on the vertical axis. The range is 0.015 in the bottom charts versus 0.3 in the top charts, so 20 times smaller for the bottom charts.

From the bottom charts we can see that the MSE ratio is slightly higher for the time-varying parameter case. But the difference is tiny. Across the five currencies, it is, on average, equal to 0.002 for $L = 120$.

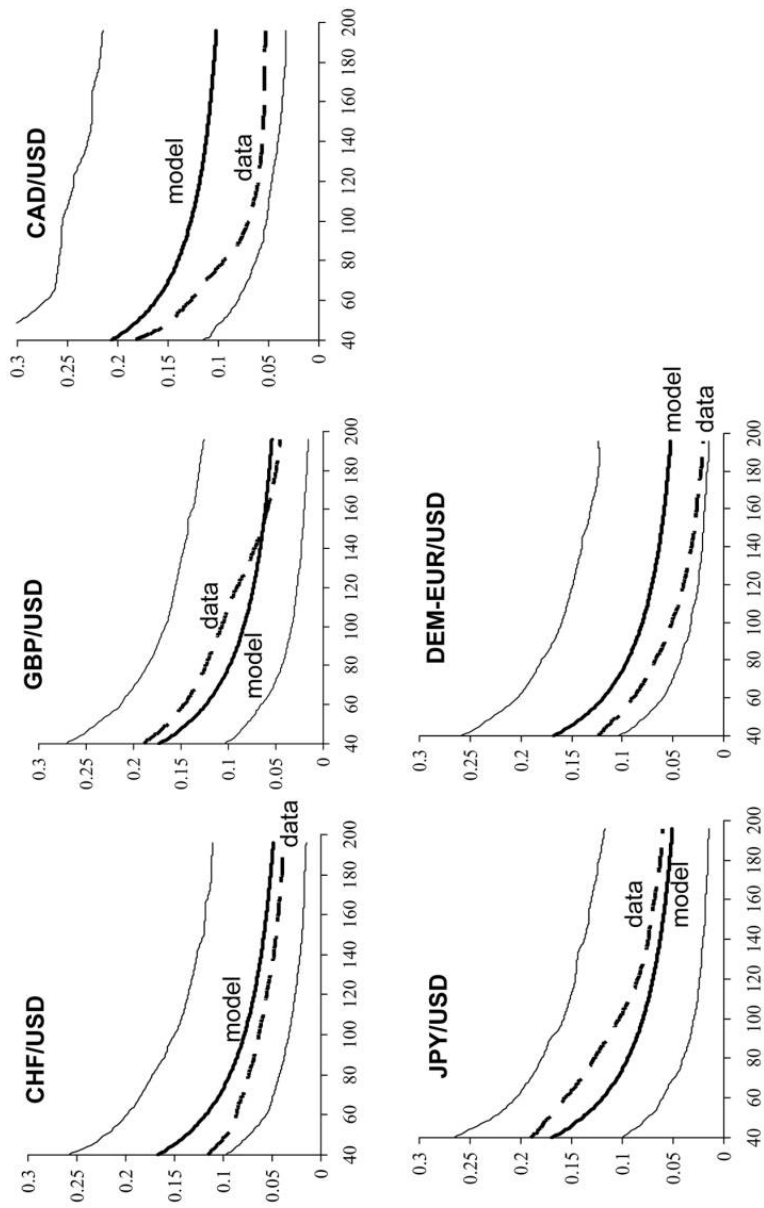
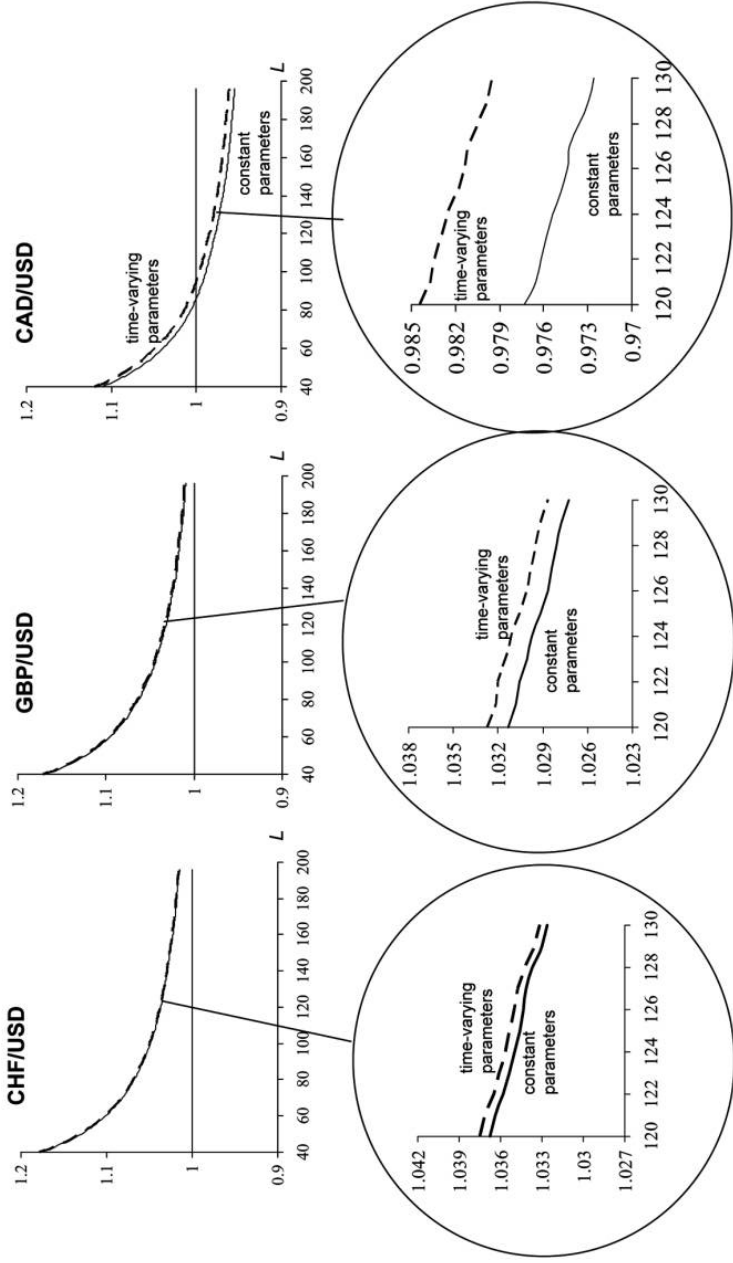


Fig. 6. In-sample fit model versus data: R^2 . Each chart contains four lines. The dotted line is the average R^2 in the data based on $P = 200$ rolling regressions of length L . The bold line represents the corresponding statistic in the model. It is an average over 1,000 simulations of the model with the same L and P as in the data. The top and bottom thin lines represent the 99% confidence interval based on 1,000 simulations of the model.



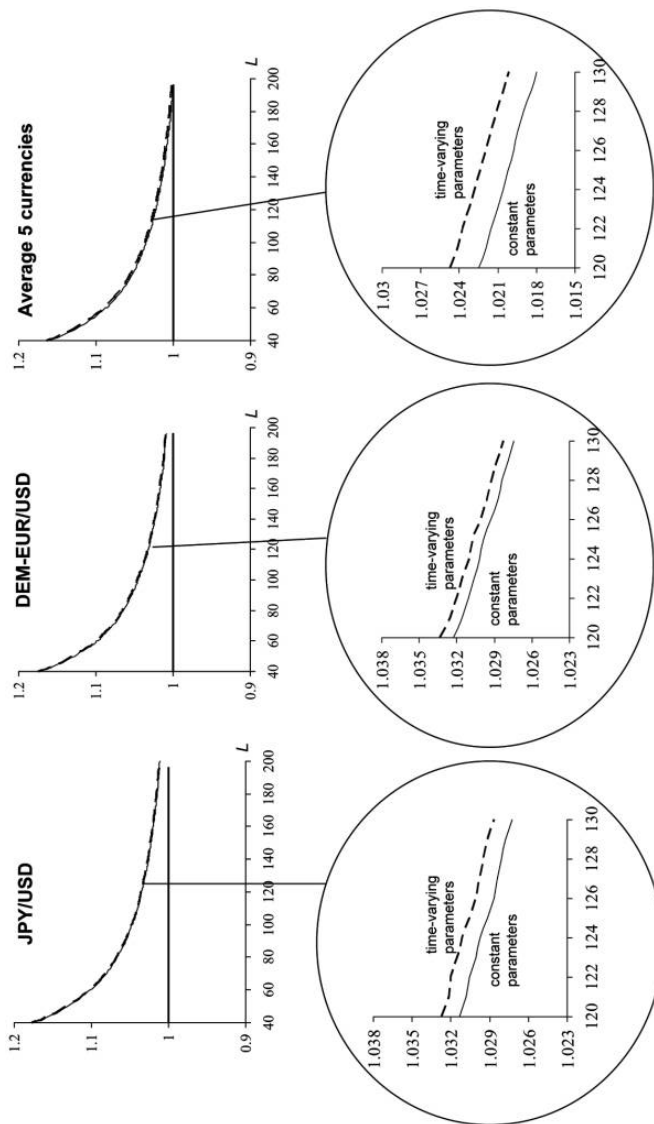
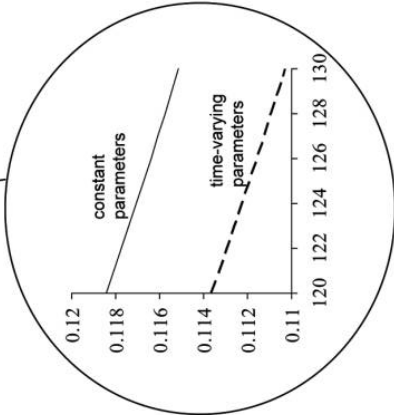

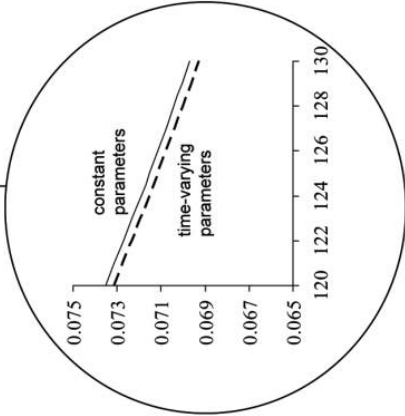





Fig. 7. Out-of-sample fit in data: MSE^{MODEL}/MSE^{RW} ($\rho_B = 0$). Each chart contains two lines. The solid line represents the ratio of the MSE of the model relative to the MSE of the random walk for the constant parameter case. The dotted line represents the same ratio for the time-varying parameter case. In the latter the persistence of the fundamentals is set at zero, and the standard deviation of parameter innovations is set at the mean level of parameters of one. MSEs are based on $P = 200$ overlapping regressions of sample length L (horizontal axis). The numbers reported represent the average over five currencies based on 1,000 simulations of the model for each of the currencies. The bottom charts zoom in on the range of sample lengths L from 120 to 130. The total length of the horizontal axis is only 0.015 in the bottom charts, versus 0.3 in the top charts (20 times bigger).



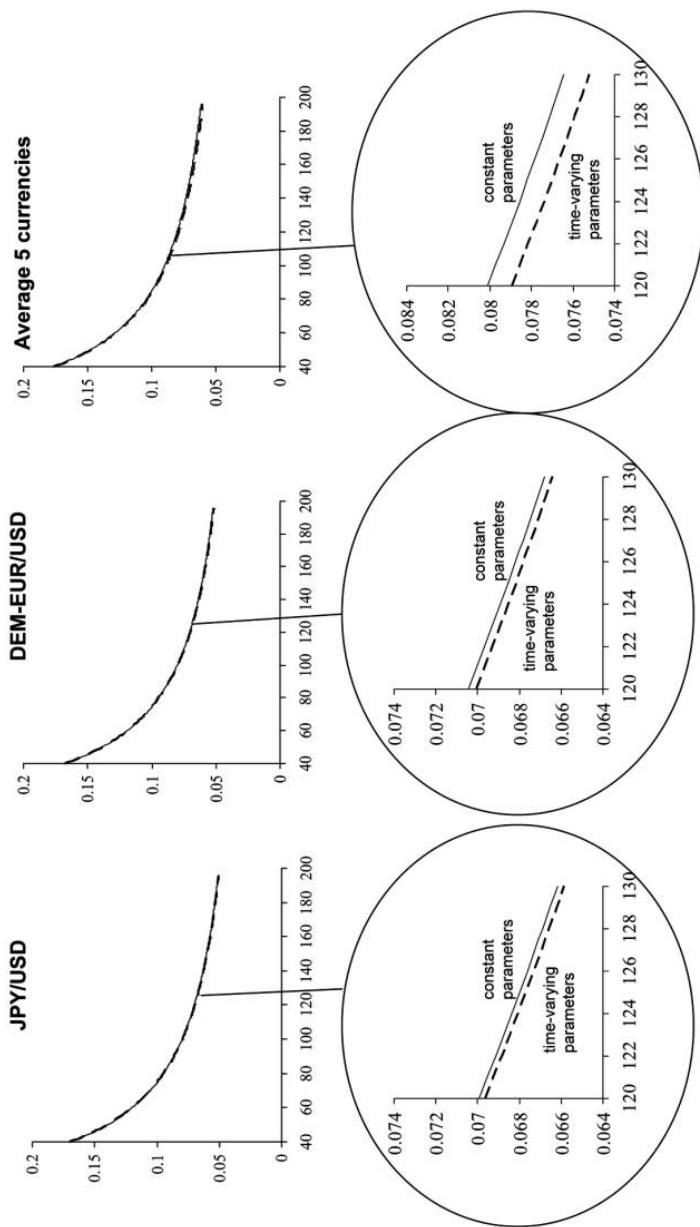


Fig. 8. In-sample fit in model ($\rho_\beta = 0$). Each chart contains two lines. The solid line represents the average R^2 for the constant parameter case. The dotted line represents the average R^2 for the time-varying parameter case. In the latter the persistence of the fundamentals is set at zero, and the standard deviation of parameter innovations is set at their mean level of parameters of one. The R^2 is an average over $P = 200$ rolling regressions of sample length L (horizontal axis). The numbers reported represent the average over five currencies based on 1,000 simulations of the model for each of the currencies. The bottom charts zoom in on the range of sample lengths L from 120 to 130. The total length of the horizontal axis is only 0.01 in the bottom charts, versus 0.2 in the top charts (20 times bigger).

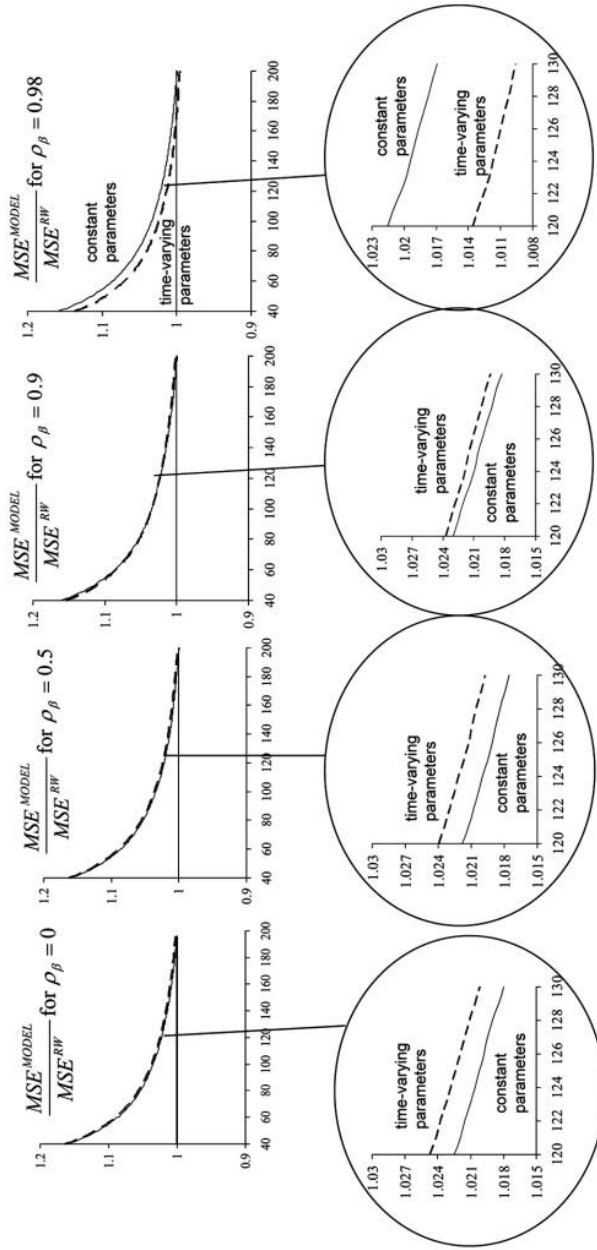


Fig. 9. Out-of-sample fit in model—sensitivity to ρ_β . Each chart contains two lines. The solid line represents the ratio of the MSE of the model relative to the MSE of the random walk for the constant parameter case. The dotted line represents the same ratio for the time-varying parameter case. In the latter the unconditional standard deviation of parameters is set at the mean level of parameters of one. The different charts report results for different values of the persistence of the parameters: 0, 0.5, 0.9, and 0.98. MSEs are based on $P = 200$ overlapping regressions of sample length L (horizontal axis). The numbers reported represent the average over five currencies based on 1,000 simulations of the model for each of the currencies. The bottom charts zoom in on the range of sample lengths L from 120 to 130. The total length of the horizontal axis is only 0.015 in the bottom charts, versus 0.3 in the top charts (20 times bigger).

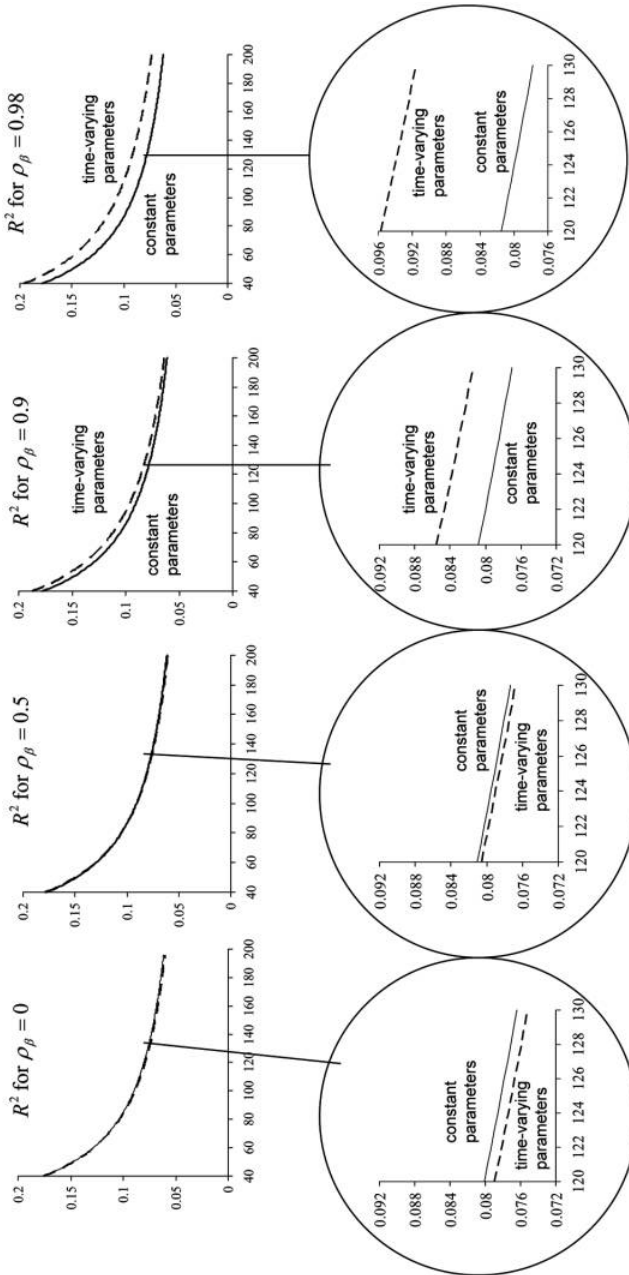


Fig. 10. In-sample fit in model—sensitivity to ρ_β . Each chart represents the average R^2 of the model in the constant parameter case. The dotted line represents the average R^2 for the time-varying parameter case. In the latter the persistence of the fundamentals is set at zero, and the unconditional standard deviation of parameters is set at the mean level of parameters of one. The R^2 is an average over $P = 200$ overlapping regressions of sample length L (which is on the horizontal axis). The numbers reported represent the average over five currencies based on 1,000 simulations of the model for each of the currencies. The bottom charts zoom in on the range of sample lengths L from 120 to 130. The total length of the horizontal axis is only 0.02 in the bottom charts, versus 0.2 in the top charts (10 times bigger).

Even for Canada, for which the difference is by far the largest and most visible by the naked eye on the top chart, it is only 0.007 for $L = 120$. While this goes in the right direction in terms of explaining the high MSE ratio, it does not amount to much quantitatively. The top chart for the average of the five currencies pretty much sums this up.

Analogous results are reported in figure 8 for the in-sample fit, the average R^2 . It is again the case that with the exception of the Canadian dollar, the difference between the average R^2 for the constant parameter case is virtually indistinguishable from that for the time-varying parameter case. The bottom charts again zoom in on the range of L from 120 to 130 with a total range of the vertical axis of 0.01 (again 20 times smaller than the top charts). It shows that the average R^2 is slightly lower for the time-varying parameter case. This is consistent with the out-of-sample fit also being slightly weaker in the time-varying parameter case. But the difference is again tiny. For the average of the five currencies, the average R^2 is 0.001 lower for the time-varying parameter case than for the constant parameter case when $L = 120$.

So far we have considered only the case in which $\rho_\beta = 0$. In figure 9 we compare the results for four different values of ρ_β : 0, 0.5, 0.9, and 0.98. In order to save space we now report only the average across the five currencies across 1,000 simulations of the model. The charts at the bottom again zoom in on the range of L from 120 to 130. These charts again have a range on the vertical axis that is 20 times narrower than for the top charts. The results for $\rho_\beta = 0.5$ are virtually identical to those for $\rho_\beta = 0$. If we increase ρ_β even further, to 0.9, the difference between the constant and time-varying parameter cases becomes even smaller. The difference is now only 0.0007 for $L = 120$.

When we increase ρ_β even further, the MSE ratio at some point becomes lower for the time-varying parameter case than for the constant parameter case. This is illustrated in the last chart of figure 9, where $\rho_\beta = 0.98$, which is close to a random walk for the parameters. In that case not only is the MSE ratio lower only for the time-varying parameter case, but the difference is not insignificant. The difference in the MSE ratio is 0.008 for $L = 120$ and an even bigger 0.013 for $L = 80$. While these numbers are not negligible, they do not help in explaining the Meese-Rogoff puzzle of underperformance of the model relative to the random walk. If anything, this makes the puzzle only worse as time-varying parameters improve the out-of-sample performance of the model relative to the random walk. We should also emphasize that this is a rather extreme case that is relevant only when ρ_β is close to one (random walk).

Otherwise the out-of-sample fit is virtually the same in the constant and time-varying parameter cases.

Figure 10 reports analogous results for the in-sample fit. Here we see a similar pattern. As we raise ρ_β , the difference between the time-varying and constant parameter cases at first becomes smaller and then changes sign. In this case the difference becomes smaller when we raise ρ_β from 0 to 0.5. The average R^2 is now 0.0005 lower for the time-varying case than for the constant parameter case for $L = 120$. When we raise ρ_β further to 0.9, the average R^2 is now higher in the time-varying parameter case. The difference is 0.005 for $L = 120$. When we raise ρ_β further to 0.98, the difference rises further to a substantial 0.014 for $L = 120$ and 0.016 for $L = 80$. This is consistent with the out-of-sample performance, where the fit is also substantially better for the time-varying parameter case when ρ_β gets close to one.

In summary, the impact of time-varying parameters on both the in- and out-of-sample fit either is close to zero or, in the case in which ρ_β is close to one, goes in the wrong direction in that it lowers the MSE ratio. We can conclude that time-varying parameters do not help to explain the weak out-of-sample performance of the model relative to the random walk.

V. What Explains the Out-of-Sample Results?

In order to shed light on the out-of-sample results of the previous section, we now derive an explicit theoretical expression for the MSE ratio in the context of a somewhat simplified version of the model. For the purpose of this theoretical exercise we simplify the model in three ways. First, we assume that the autoregressive coefficients on the fundamentals are the same for all fundamentals. Second, we assume that $\Omega_f = \sigma_f^2 I_N$, where I_N is the identity matrix of size N . This is therefore a symmetric case in which all fundamentals have the same standard deviation and their innovations have zero correlation. Finally, we assume that $\rho_u = 0$, which in any case is close to our calibration results.

In addition to these simplifications, we will compute the MSE ratio only for the case in which $P = \infty$. In other words, we will consider an infinite number of rolling regressions over samples of length L . Then $\text{MSE}^{\text{MODEL}}$ and MSE^{RW} will be equal to the expectation of the MSEs for any particular sample of length L . Moreover, we compute the MSE of the model without a constant term. While these changes will not make the results completely comparable to those in Section IV, they will nonetheless shed clear insight into what factors drive the out-of-sample

forecasting performance of the model relative to the random walk and particularly the role of time-varying parameters.

We first compute the expected MSE for the random walk prediction. This is equal to the unconditional expectation of Δs_t^2 . Since

$$\Delta s_t = \sum_{n=1}^N \beta_{nt} f_{nt} + u_t, \quad (7)$$

our assumptions above imply that

$$\text{MSE}^{\text{RW}} = E(\Delta s_t^2) = N \text{Var}(f) [\beta^2 + \text{Var}(\beta)] + \text{Var}(u), \quad (8)$$

where

$$\text{Var}(\beta) = \frac{\sigma_\beta^2}{1 - (\rho_\beta)^2}, \quad \text{Var}(u) = \sigma_u^2, \quad \text{Var}(f) = \frac{\sigma_f^2}{1 - (\rho_f)^2}.$$

The expected MSE of the model is a bit more complicated to compute. We describe the main results, leaving details of the algebra to appendix B. The first step is to estimate the parameters from a regression of Δs_t on the fundamentals over a sample of L periods. Assume that the regression uses data from $t - L$ to $t - 1$ and the results are used at $t - 1$ to forecast Δs_t . Let δ denote the vector of estimated parameters of the N fundamentals. Define $f_t = (f_{1t}, \dots, f_{Nt})'$. We have

$$\hat{\delta} = \left(\sum_{i=1}^L f_{t-i} f_{t-i}' \right)^{-1} \sum_{i=1}^L f_{t-i} \Delta s_{t-i} = \sum_{i=1}^L \lambda_i \beta_{t-i} + \psi_t \sum_{i=1}^L u_{t-i} f_{t-i}, \quad (9)$$

where the weights λ_i are matrices that sum to the identity matrix I :

$$\lambda_i = \left(\sum_{j=1}^L f_{t-j} f_{t-j}' \right)^{-1} f_{t-i} f_{t-i}'$$

and

$$\psi_t = \left(\sum_{j=1}^L f_{t-j} f_{t-j}' \right)^{-1}$$

is a matrix as well. The estimate $\hat{\delta}$ therefore has two components. The first is a weighted average of past parameters β_{t-i} . The second is a component due to small-sample estimation error. This last component will vanish to zero when the sample length L approaches infinity.

Using expression (9) for the estimated parameters in the rolling regressions, we have

$$\begin{aligned} \text{MSE}^{\text{MODEL}} &= E(\Delta s_t - f_t' \hat{\delta})^2 \\ &= \text{Var}(f) E \left(\beta_t - \sum_{i=1}^L \lambda_i \beta_{t-i} \right)' \left(\beta_t - \sum_{i=1}^L \lambda_i \beta_{t-i} \right) \\ &\quad + \text{Var}(u) E(f_t' \psi_t f_t) + \text{Var}(u). \end{aligned} \quad (10)$$

Apart from the noise shocks u_t , two factors drive the MSE in the model. Both are related to the fact that the future parameter is unknown and needs to be estimated. The first is the standard small-sample estimation error. This is captured by the term $\text{Var}(u)E(f_t'\psi_t f_t)$ and applies equally under constant and time-varying parameters. Second, in the presence of time-varying parameters, there is an additional source of estimation error. Even when we abstract from small-sample estimation error, the parameter estimate is a weighted average of parameters over the past L periods. This weighted average of past parameters will differ from the parameter vector tomorrow when parameters change over time. This is captured by the term

$$\text{Var}(f) E \left(\beta_t - \sum_{i=1}^L \lambda_i \beta_{t-i} \right)' \left(\beta_t - \sum_{i=1}^L \lambda_i \beta_{t-i} \right).$$

Both of these sources of parameter estimation error raise the MSE.

In appendix B, we derive an expression for the estimation error that is due to time-varying parameters:

$$E \left(\beta_t - \sum_{i=1}^L \lambda_i \beta_{t-i} \right)' \left(\beta_t - \sum_{i=1}^L \lambda_i \beta_{t-i} \right) = N \text{Var}(\beta) h, \quad (11)$$

where

$$h = z \frac{2}{L-1} \sum_{i=1}^{L-1} (L-i) [1 - (\rho_\beta)^i] + \frac{2}{L(L-1)} \sum_{i=2}^L (i-1) [1 - (\rho_\beta)^i],$$

and $z > 0$ is the expectation of any element on the diagonal of $\lambda_i' \lambda_i$ (for any i).

We can now evaluate the implications of time-varying parameters for the MSE ratio. We have

$$\begin{aligned} \frac{\text{MSE}^{\text{MODEL}}}{\text{MSE}^{\text{RW}}} &= 1 + \frac{\text{MSE}^{\text{MODEL}} - \text{MSE}^{\text{RW}}}{\text{MSE}^{\text{RW}}} \\ &= 1 + \frac{-N \text{Var}(f)\beta^2 + \text{Var}(u)E(f'_t \psi_t f_t)}{N \text{Var}(f)[\beta^2 + \text{Var}(\beta)] + \text{Var}(u)} \\ &\quad + \frac{N \text{Var}(f) \text{Var}(\beta)(h-1)}{N \text{Var}(f)[\beta^2 + \text{Var}(\beta)] + \text{Var}(u)}. \end{aligned} \quad (12)$$

First set $\text{Var}(\beta) = 0$, so that parameters are constant. Then the last fraction is zero. We will refer to the first fraction as M_c , which is the MSE ratio minus one under constant parameters:

$$M_c = \frac{-N \text{Var}(f)\beta^2 + \text{Var}(u)E(f'_t \psi_t f_t)}{N \text{Var}(f)\beta^2 + \text{Var}(u)}. \quad (13)$$

It has two parts. First, to the extent that fundamentals have explanatory power, the model's performance is better than the random walk. This is captured by the term $-N \text{Var}(f)\beta^2$ in the numerator. Second, small-sample estimation error of parameters deteriorates the out-of-sample performance of the model relative to the random walk. This is captured by the term $\text{Var}(u)E(f'_t \psi_t f_t) > 0$ in the numerator. This latter effect tends to dominate, especially for small samples L , so that $M_c > 0$ and the MSE ratio is larger than one.

Now consider the impact of time-varying parameters. The impact comes through three channels. First, it raises the MSE for the random walk as Δs_t becomes a bit more volatile because of time-varying parameters. This by itself reduces the MSE ratio, assuming that $M_c > 0$. It is reflected by the increase in $\text{Var}(\beta)$ in the denominator in the first ratio of (12). Second, time-varying parameters raise the estimation error of the future parameter. The estimation is now of a weighted average of past parameters, which is not equal to the future parameter. This additional estimation error, which comes on top of the small-sample estimation error that equally applies under constant parameters, is captured by $h > 0$ in the second ratio of (12). This deteriorates the out-of-sample performance of the model relative to the random walk. Third, abstracting from estimation error, time-varying parameters increase the explanatory power of fundamentals as they raise the expectation of the squared parameters. This lowers the MSE ratio and is captured by the term -1 after the h in the numerator of the last ratio in (12).

This last point is perhaps most clearly illustrated by considering a case of time-varying parameters in which the parameters are known, so that we can completely abstract from estimation error. The variance of the component of Δs_t that is explained by fundamentals is then

$$\text{Var}\left(\sum_{n=1}^N \beta_{nt} f_{nt}\right) = N \text{Var}(f) E\beta_{nt}^2 = N \text{Var}(f)[\beta^2 + \text{Var}(\beta)], \quad (14)$$

which rises with parameter volatility. Equation (14) shows that what matters for the explanatory power of fundamentals is not the mean level of parameters, but the expectation squared level of parameters, $E\beta_{nt}^2$, which rises with parameter volatility.

The increased MSE for the random walk and the increased explanatory power of fundamentals under time-varying parameters (increase in $E\beta_{nt}^2$) both reduce the MSE ratio. On the other hand, the increased estimation error of the parameters raises the MSE ratio. We would like to know how this adds up and what the effect is quantitatively. In order to do so, we will take the derivative of the MSE ratio in (12) with respect to $\text{Var}(\beta)$ at the point at which $\text{Var}(\beta) = 0$ (constant parameter case). Also setting $\beta = 1$, we get

$$\Delta \left(\frac{\text{MSE}^{\text{MODEL}}}{\text{MSE}^{\text{RW}}} \right) = R_c^2 (-M_c + h - 1) d \text{Var}(\beta), \quad (15)$$

where

$$R_c^2 = \frac{N \text{Var}(f)}{N \text{Var}(f) + \text{Var}(u)}$$

is the infinite sample R^2 in the constant parameter case.

The effect of time-varying parameters depends on both R_c^2 and the term $-M_c + h - 1$ that reflects the increase in MSE of the random walk ($-M_c$), the increased estimation error of the parameters ($+h$), and the increased explanatory power of the fundamentals under time-varying parameters (-1). With respect to the R^2 , for the average of the five currencies we have $R_c^2 = 0.032$. Clearly, the quantitative effect of time-varying parameters is reduced by the fact that the explanatory power of the fundamentals is quite limited. The parameters do not matter much if the fundamentals that multiply them do not have much explanatory power for Δs_t in the first place.

Equation (15) allows us to break down the impact of time-varying parameters into three components: the effect of the increase in the MSE of the random walk, the increased estimation error of the parameters, and the increased explanatory power of the fundamentals. We provide this breakdown in figures 11 and 12. In doing so we set the variance of the fundamentals equal to the average across the five fundamentals from Section IV, and we set σ_u to match the average standard deviation of the exchange rate change across the five currencies. We also set $\text{Var}(\beta) = 1$ as we did in Section IV.

Figure 11 shows results for the same four different values of ρ_β considered in Section IV: 0, 0.5, 0.9, and 0.98. Each chart shows the total impact of time-varying parameters on the MSE ratio (change relative to the constant parameter case) as well as the role of each of the three contributing factors. We should point out that while the exercise is not exactly comparable to that in Section IV because of various simplifications that we adopted in this section, the total impact of time-varying parameters on the MSE ratio is nonetheless very close to that reported in Section IV and shown in figure 9.

Several points can be made from figure 11. First, as before, the total impact of time-varying parameters is tiny. The only exception is again the case of $\rho_\beta = 0.98$, when time-varying parameters reduce the MSE ratio by an amount that is nontrivial for low L . Second, in terms of the breakdown the largest impact comes from the rise in the MSE ratio due to increased estimation error and the drop in the MSE ratio due to increased explanatory power of the fundamentals. But these two factors almost exactly offset each other. The only exception is again the case in which $\rho_\beta = 0.98$. In that case the increase in the MSE ratio due to increased estimation error of the parameters is dominated by the increased explanatory power of the fundamentals. When $\rho_\beta \rightarrow 1$, then $h \rightarrow 0$ and time-varying parameters do not lead to increased estimation error of the parameters. The reason is that parameters become highly persistent and the standard deviation of parameter innovations goes to zero when $\rho_\beta \rightarrow 1$ and we hold $\text{Var}(\beta)$ constant.

The role of ρ_β is further illustrated in figure 12. It is analogous to figure 11 except that we keep $L = 120$ and now vary ρ_β from zero to one. Figure 12 clearly shows that ρ_β plays a role only when it gets very close to one. In that case the increased estimation error goes away and the increased explanatory power of the fundamentals leads to a substantial reduction in the MSE ratio. But unless ρ_β is nearly one, we can conclude that the impact of time-varying parameters on the MSE ratio is negligible. The increase in the MSE ratio due to increased estimation error of

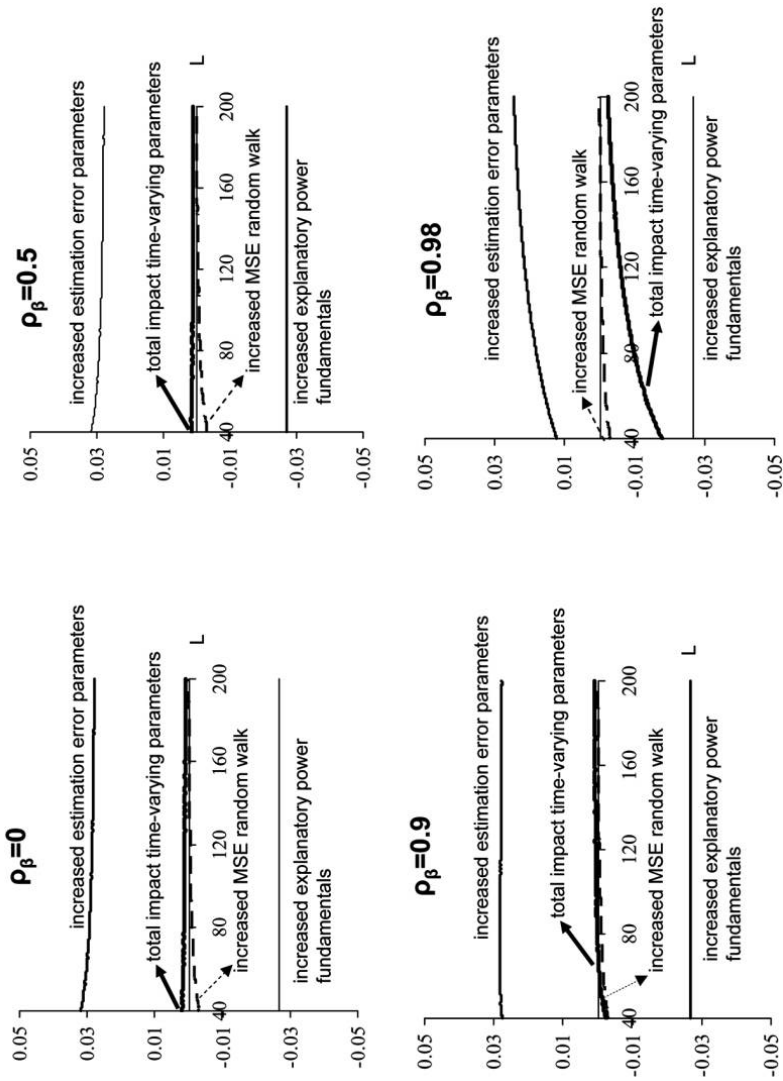


Fig. 11. Impact of time-varying parameters on out-of-sample fit




Fig. 12. Impact of time-varying parameters on out-of-sample fit: role of ρ_β

the parameters is almost exactly offset by the decrease due to the increased explanatory power of the fundamentals. Even if we believe that ρ_β is very close to one, it would not help to resolve the Meese-Rogoff puzzle: instead of explaining the poor performance of the model, time-varying parameters improve it relative to the random walk in that case.

Since time-varying parameters have little overall impact on the MSE ratio, we conclude that the underperformance of the model relative to the random walk is entirely due to small-sample estimation error of parameters that applies equally under time-varying and constant parameters.¹² Figure 13 illustrates this for the case of constant parameters. As can be seen from the expression for M_c in (13), two factors contribute to the MSE ratio under constant parameters. First, if the parameters were known, the model clearly outperforms the random walk, lowering the MSE ratio by about 0.03. Second, small-sample estimation error of the parameters raises the MSE ratio. It is this second factor that dominates, especially for low L , and explains why the model does not outperform the random walk. Additionally introducing time-varying parameters does not significantly change this conclusion.

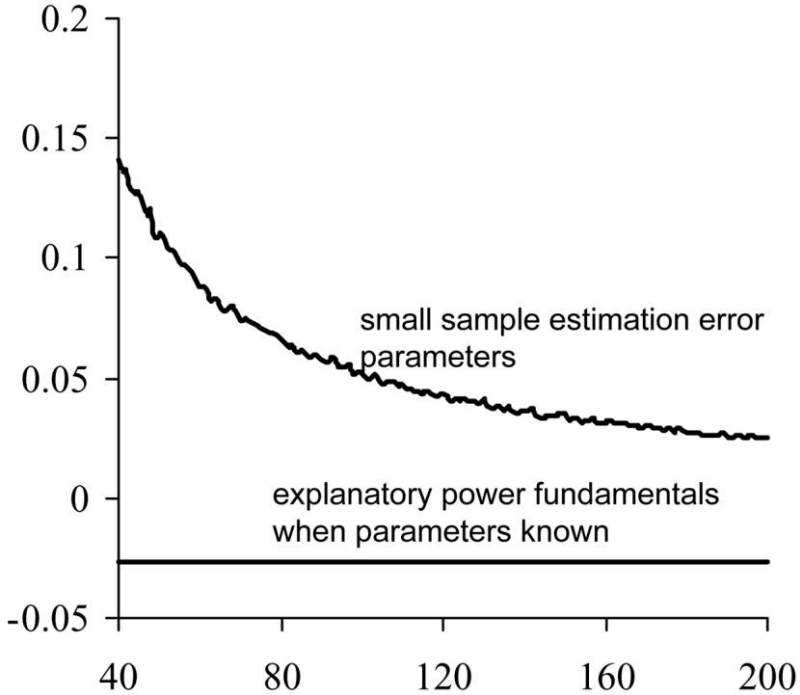


Fig. 13. Factors contributing to out-of-sample fit under constant parameters

VI. What Explains the In-Sample Results?

In the previous section we looked at what factors drive the impact of time-varying parameters on the out-of-sample fit. In this section we do the same for the in-sample fit. Define $\hat{\delta}_t$ as the estimated vector of coefficients in a regression of the change in the log exchange rate on the fundamentals over the sample $(t - L, t - 1)$. This is

$$\begin{aligned} \hat{\delta}_t &= \left(\sum_{i=1}^L f_{t-i} f'_{t-i} \right)^{-1} \sum_{i=1}^L f_{t-i} \Delta s_{t-i} \\ &= \sum_{i=1}^L \lambda_i \beta_{t-i} + \psi_t \sum_{i=1}^L u_{t-i} f_{t-i}. \end{aligned} \tag{16}$$

The R^2 of this regression is

$$R^2 = 1 - \frac{(1/L) \sum_{j=1}^L (\Delta s_{t-j} - f'_{t-j} \hat{\delta}_t)^2}{(1/L) \sum_{j=1}^L \Delta s_{t-j}^2}. \tag{17}$$

So far in both the data and the model we have defined the in-sample fit as the average R^2 over the $P = 200$ rolling regressions. For the purpose of this section we define it slightly differently. We define it as in (17) but with an expectation in both the numerator and denominator. One can think of the ratio in (17) as a ratio of two MSEs: the in-sample mean squared regression error and the mean squared random walk forecast error. The average R^2 is equal to one minus the average ratio of these mean squared errors across the rolling regressions. Instead we now define the in-sample fit as one minus the ratio of the average of these mean squared errors across $P = \infty$ rolling regressions. This significantly facilitates the analysis, and in both the data and model simulations, it makes little difference whether one uses the average ratio or the ratio of the average for the in-sample fit.

Defining the in-sample fit as in (17) with the expectation in both the numerator and denominator, after some algebra that is in appendix C we obtain

$$R^2 = \frac{N \text{Var}(f)\beta^2 + (N/L) \text{Var}(u)}{N \text{Var}(f)[\beta^2 + \text{Var}(\beta)] + \text{Var}(u)} + \frac{N \text{Var}(f) \text{Var}(\beta)(1 - \omega)}{N \text{Var}(f)[\beta^2 + \text{Var}(\beta)] + \text{Var}(u)}, \quad (18)$$

where

$$\omega = 2 \left[\frac{L-2}{L^2(L-1)} + \frac{1}{L-1} z \right] \sum_{i=1}^{L-1} (L-i)[1 - (\rho_\beta)^i]. \quad (19)$$

The discussion of what drives this in-sample R^2 will parallel that for the case of the out-of-sample MSE ratio. First consider the case of constant parameters, so that $\text{Var}(\beta) = 0$. Then the in-sample fit as a function of L becomes

$$R_c^2(L) = \frac{N \text{Var}(f)\beta^2 + (N/L) \text{Var}(u)}{N \text{Var}(f)\beta^2 + \text{Var}(u)}. \quad (20)$$

Two factors drive the in-sample fit. First, the explanatory power of the fundamentals that is reflected in the term $N \text{Var}(f)\beta^2$ in the numerator raises the R^2 . Second, even in the absence of any explanatory power of the fundamentals ($\beta = 0$), the R^2 is still positive because of a spurious small-sample fit. This is captured by the second term in the numerator, $(N/L) \text{Var}(u)$. This term goes to zero as the sample length L goes to infinity. So while small-sample estimation error of the parameters deteriorates the out-of-sample fit, it improves the in-sample fit.

We now consider the impact of time-varying parameters. The impact comes through three channels. First, it raises the variance of Δs_t , which by itself (holding constant the explanatory power of the fundamentals) lowers the R^2 . This is captured by $\text{Var}(\beta)$ in the denominator of the first ratio in (18). Second, time-varying parameters raise the estimation error of parameters. This is captured by $\omega > 0$:

$$\frac{1}{L} \sum_{j=1}^L E \left(\beta_{t-j} - \sum_{i=1}^L \lambda_i \beta_{t-i} \right)' \left(\beta_{t-j} - \sum_{i=1}^L \lambda_i \beta_{t-i} \right) = N \text{Var}(\beta) \omega. \quad (21)$$

Estimation over the period $t - L$ to $t - 1$ leads to an estimate of a weighted average of parameters, $\sum_{i=1}^L \lambda_i \beta_{t-i}$. Even when we abstract from small-sample estimation error, this weighted average of past parameters differs from the actual parameters β_{t-j} that vary over time. This lowers the in-sample fit as captured by the term $-\omega$ in the numerator of the second ratio in (18). Finally, time-varying parameters raise the explanatory power of the fundamentals as they raise the expected squared value of the parameters. This improves the R^2 and is reflected by the 1 before the $-\omega$ in the numerator of the second ratio in (18).

While small-sample estimation error has an opposite impact on the in- and out-of-sample fit, time-varying parameters have a very similar effect. The three factors related to time-varying parameters are analogous to those discussed for the out-of-sample fit: increased MSE of the random walk (increased variance of Δs_t), increased estimation error of parameters, and increased explanatory power of fundamentals due to a rise in $E\beta_{nt}^2$. The last factor improves both the in- and out-of-sample fit. The second factor deteriorates both the in- and out-of-sample fit. Only the first factor operates in opposing directions, improving the out-of-sample fit while deteriorating the in-sample fit, but it is tiny compared to the other two factors.

Analogously to the discussion of the out-of-sample fit, we can again evaluate the quantitative effect of time-varying parameters by differentiating (18) with respect to $\text{Var}(\beta)$ at the point at which $\text{Var}(\beta) = 0$. Setting $\beta = 1$, we have

$$dR^2 = R_c^2 [-R_c^2(L) + 1 - \omega] d \text{Var}(\beta). \quad (22)$$

One can again expect the effect to be small to the extent that the explanatory power of the fundamentals, captured by R_c^2 , is small. The three factors driving the impact of time-varying parameters on the out-of-sample fit are shown in brackets: the $-R_c^2(L)$ term that captures the increased variance of Δs_t , the $+1$ term that captures the increased explanatory power of

the fundamentals, and the $-\omega$ term that captures the increased estimation error of the parameters.

Without repeating all the graphics analogous to figures 11 and 12, the message is the same. The second factor (increased estimation error parameters) lowers the in-sample fit by almost the same amount as the third factor (increased explanatory power fundamentals) raises it, whereas the first factor is quite small. It is again the case that only when ρ_β is very close to one does the increased estimation error disappear. Only in that case is there a significant effect of time-varying parameters, which raises the in-sample fit just as it significantly improves the out-of-sample fit in that case.

Since with the exception of $\rho_\beta \rightarrow 1$ the impact of time-varying parameters is small, we conclude that the in-sample R^2 is mainly driven by the two factors relevant under constant parameters: the true explanatory power of the fundamentals and the small-sample estimation error. Both raise the R^2 , with the small-sample estimation error clearly dominant over small samples. This is illustrated in figure 14.




Fig. 14. Factors contributing to in-sample fit (R^2) under constant parameters

VII. Additional Results

In this section we provide some additional results that largely confirm our findings so far. First, in figure 15 we provide some additional sensitivity analysis results. Figure 15 has four charts. Each chart shows the MSE ratio for the out-of-sample forecast, with parameter values corresponding to Japan, based on an average of 1,000 simulations of the model. The results under both constant (solid line) and time-varying parameters (dotted line) are shown. In the case of time-varying parameters it is assumed that parameter innovations have no persistence, so $\rho_\beta = 0$.

The top two charts show results when we forecast 3 and 12 months out of sample rather than the 1-month forecast considered so far. Again the difference between the cases of constant and time-varying parameters is virtually nil. The two lines are almost exactly on top of each other. The bottom-left chart considers a Markov process instead of a process with normally distributed parameter innovations that we have considered so far. There are two states, in which the parameter takes on the values of, respectively, zero and two, both with equal probability. Given any state we are in, there is an equal probability of 0.5 of staying in that state or moving to the other state. This process implies as before that both the mean and standard deviation of the parameters are one and that they are uncorrelated over time. It is clear from figure 15 that this alternative process makes no difference for the results. The constant and time-varying parameter cases are again virtually indistinguishable.

We have also simulated the model under a Markov process in which transition probabilities are different from 0.5. We assume that the probability of the parameter staying in the same state is p and the probability of transitioning to the other state is $1 - p$, with the two states again being zero and two. In that case the first-order autocorrelation of the parameter is $2p - 1$. Consistent with the results for the AR process, we find that the constant and time-varying parameter cases remain indistinguishable unless the autocorrelation is close to one (p close to one). In the latter case we again find that time-varying parameters lower the MSE ratio and therefore cannot explain the Meese-Rogoff puzzle of a high MSE ratio.

Finally, the bottom-right chart shows the result when we triple the standard deviation of fundamental innovations. This leads to an MSE ratio well below one for both the constant and time-varying parameter cases since the fundamentals have significantly more explanatory power. But now the MSE ratio is visibly higher for the time-varying

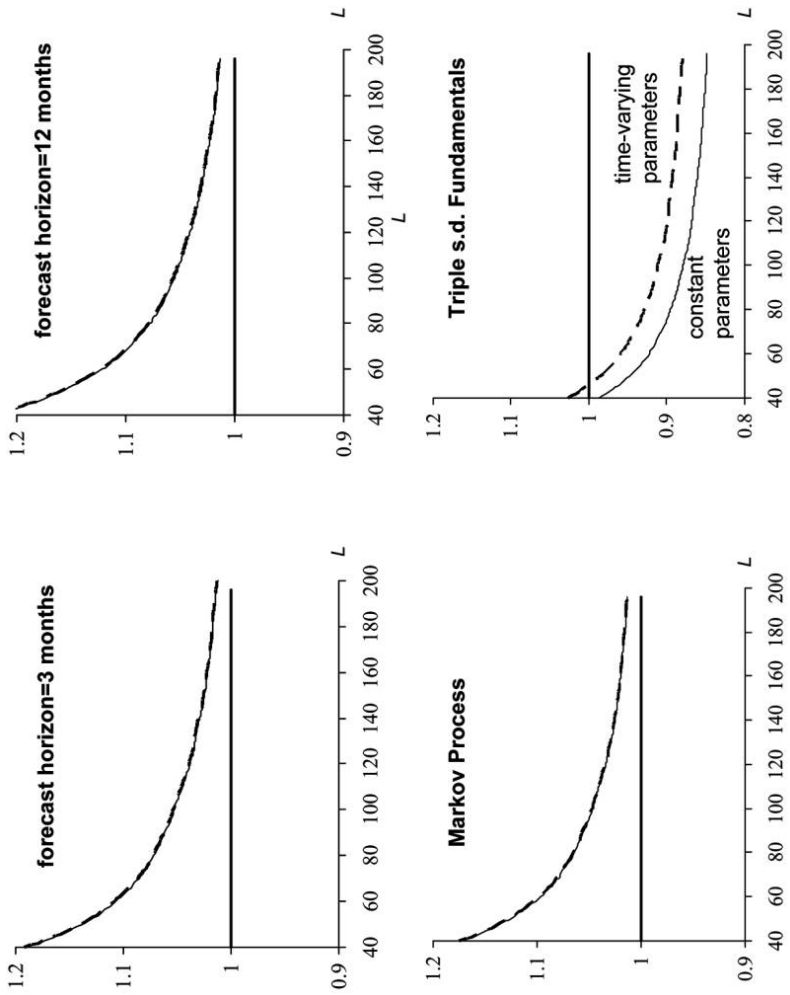


Fig. 15. Sensitivity analysis—Japan

parameter case. This is not surprising in light of the findings of the previous section and can be understood directly from (15). The explanatory power of fundamentals as measured by R_c^2 is now multiplied by almost a factor 9 as $\text{Var}(f)$ is multiplied by a factor 9. The higher MSE ratio under time-varying parameters, which previously was visible only after significantly zooming in on the numbers (see fig. 7), now becomes substantial. This shows that when fundamentals have significant explanatory power, time-varying parameters can make a difference for the MSE ratio.

It is unusual to find empirical exchange rate equations with a high R^2 . One exception is the case of commodity currencies presented in Section II and figure 6. In such a case one might indeed argue that the MSE ratio would have been even lower without time-varying parameters. But in this case, there is no Meese-Rogoff puzzle to solve since the model beats the random walk.

Finally, figure 16 shows from a somewhat different perspective that the difference between constant and time-varying parameters is small. Figure 3 showed the evolution of estimates of a particular parameter. Figure 16 does the same using data generated by the reduced-form model. For one particular simulation of the model we report the estimated parameter coefficient for variable 1 for each of the $P = 200$ rolling regressions. The horizontal axis shows the number of the rolling regression. The results are reported for regressions of length $L = 40$, $L = 120$, and $L = 200$. Each chart shows the result for both constant parameters (thick line) and time-varying parameters (thin line) with $\rho_\beta = 0$. Clearly, the estimated parameter varies significantly across rolling regressions. It varies from about -10 to $+8$ for $L = 40$. As expected, it varies less when $L = 120$ and even less when $L = 200$. But even for $L = 200$ the estimated parameter varies over a range of about three.

Two points stand out. First, the variation in estimated parameters is very similar to that in the data reported in figure 3. Second, the time variation in the parameter estimates is entirely the result of small-sample estimation error. It makes virtually no difference whether the actual parameters are constant or time varying.

It is also noteworthy that even for $L = 200$ there is very large small-sample estimation bias. While the true parameter is one under constant parameters, the estimated parameter varies from zero to -3 . This estimation error explains why the MSE ratio generally continues to be above one in the data even for such long samples that are now available. The limited explanatory power of the fundamentals is more than offset by




Fig. 16. Estimated coefficients, rolling regressions model (variable 1)

the small-sample parameter estimation error. This is in our view the real explanation for the Meese-Rogoff puzzle, not the presence of time-varying parameters.

VIII. Conclusion

A priori, the unstable relationship between the exchange rate and fundamentals is a natural explanation for the poor out-of-sample forecasting performance of exchange rate models. Such instability increases parameter estimation errors. It implies that the relationships based on past behavior are less likely to be useful in the future. While this reasoning is correct, our analysis shows that there is another offsetting effect at work. Time-varying parameters tend to increase the explanatory power of fundamentals. We find that, on net, time-varying parameters have virtually no effect on the out-of-sample forecasting performance of exchange rate models. There are two exceptions to this, but neither sheds any light on the Meese-Rogoff puzzle. One is the case in which the persistence of parameters is close to one, but in that case time-varying parameters have an impact that operates in the wrong direction: it improves the out-of-sample fit of exchange rate models. This cannot explain the Meese-Rogoff puzzle of poor out-of-sample fit. The other case is one in which fundamentals have high explanatory power for the exchange rate. But this is counterfactual and implies that there is no Meese-Rogoff puzzle in the first place.

We conclude that the Meese-Rogoff puzzle can be explained only by short-sample problems. It is important to notice, however, that a major reason behind the results is that fundamentals have a low explanatory power in exchange rate equations. Even if we could solve the small-sample problem (by having infinitely long samples), in most cases we would not do much better than the random walk. This means that the basic problem is not so much the instability in the relationship between exchanges rates and fundamentals, but its weakness.

Appendix A

Data Appendix

In this appendix, we describe the data used in the paper. The first part relates to the exchange rate model based on five currencies and five macro fundamentals. The second part relates to the so-called commodity currencies.

A. Exchange Rates and Macro Fundamentals

Exchange rate: We use bilateral U.S. dollar end-of-period exchange rates from International Financial Statistics (IFS). The five currencies considered are the Swiss franc, the British pound, the Canadian dollar, the Japanese yen, and the German mark. Since the introduction of the euro in 1999, we convert the euro exchange rate to German marks using the fixed conversion factor (1.95583 marks per euro). The five macro fundamentals we consider are as follows.

Money supply: $\Delta(m_t - m_t^{\text{US}})$, where $m_t = \ln M_t$ and M_t is M1, OECD Main Economic Indicators (MEI), for Canada and M1, IFS line 59MA, for Japan. In the case of Germany/Euro Area, we consider M1 seasonally adjusted, IFS line 59MACZF until December 1998 and M1, OECD MEI, for the Euro Area from January 1999. For the United Kingdom, we take M0, IFS line 19MC.ZF, until April 2006 (last observation of the IFS series) and M1, OECD MEI, from May 2006. For Switzerland, we use IFS line 34ZF. Finally, for the United States, we take the corresponding series, that is, either M1, IFS line 59MA, or M1, OECD MEI. All seasonally unadjusted series were adjusted using monthly dummies.

Industrial production: $\Delta(y_t - y_t^{\text{US}})$, where $y_t = \ln Y_t$ and Y_t is the industrial production index, taken from IFS, line 66CZF, except for Switzerland, for which no monthly series is available. For this country, we compute monthly observations from quarterly data (IFS, line 66) using the same procedure as in Molodtsova and Papell (2009).

Unemployment rate: $\Delta(u_t - u_t^{\text{US}})$, where $u_t = \ln U_t$ and U_t is the unemployment rate from OECD MEI except for Germany/Euro Area. For this country, we take a series from Datastream (mnemonic WGUN% TOTQ) that covers only West Germany and is thus unaffected by the German reunification that took place in 1990.

Interest rate: $i_{t-1} - i_{t-1}^{\text{US}}$, where i_t is the monthly return calculated from the money market rate, IFS line 60B.

Oil price: Δp_t^{oil} , where $p_t^{\text{oil}} = \ln P_t^{\text{oil}}$ and P_t^{oil} is the average crude oil spot price from IFS.

B. Commodity Currencies

Exchange rate: As in the previous model, we use bilateral U.S. dollar end-of-period exchange rates from IFS. The three commodity currencies considered are the Australian, the Canadian, and the New Zealand dollars. We use monthly data over a common sample from January 1986 to December 2008.

Commodity prices: Δcp_t , where $cp_t = \ln CP_t$ and CP_t is a U.S. dollar-denominated country-specific index of commodity prices. For Australia, we use the index of commodity prices (“all items”) from the Reserve Bank of Australia. For Canada, the index (“total, all commodities”) is from the Bank of Canada and is obtained from the CANSIM database. Finally, for New Zealand, we use the ANZ Commodity Price Index.

Appendix B

Algebra Section V

In this appendix we will derive the result that the ratio of the expected mean squared forecast error of the model relative to the random walk can be written as (12). We already know that

$$\text{MSE}^{\text{RW}} = E(\Delta s_t^2) = N \text{Var}(f)[\beta^2 + \text{Var}(\beta)] + \text{Var}(u). \quad (\text{B1})$$

We therefore only need to compute the expected mean squared forecast error of the model, which is

$$\begin{aligned} \text{MSE}^{\text{MODEL}} &= E(\Delta s_t - f_t' \hat{\delta})^2 = E \left[f_t' \left(\beta_t - \sum_{i=1}^L \lambda_i \beta_{t-i} \right) \right. \\ &\quad \left. + u_t - f_t' \Psi_t \sum_{i=1}^L u_{t-i} f_{t-i} \right]^2 \\ &= E \left(\beta_t - \sum_{i=1}^L \lambda_i \beta_{t-i} \right)' f_t f_t' \left(\beta_t - \sum_{i=1}^L \lambda_i \beta_{t-i} \right) + \text{Var}(u) \\ &\quad + \sum_{i=1}^L \sum_{j=1}^L E u_{t-i} u_{t-j} f_t' \Psi_t f_{t-i} f_{t-j}' \Psi_t' f_t \\ &= \text{Var}(f) E \left(\beta_t - \sum_{i=1}^L \lambda_i \beta_{t-i} \right)' \left(\beta_t - \sum_{i=1}^L \lambda_i \beta_{t-i} \right) \\ &\quad + \text{Var}(u) + \text{Var}(u) E(f_t' \Psi_t f_t). \end{aligned} \quad (\text{B2})$$

We have

$$\begin{aligned}
 & E \left(\beta_t - \sum_{i=1}^L \lambda_i \beta_{t-i} \right)' \left(\beta_t - \sum_{i=1}^L \lambda_i \beta_{t-i} \right) \\
 &= N \text{Var}(\beta) + E \left(\sum_{i=1}^L \lambda_i \tilde{\beta}_{t-i} \right)' \left(\sum_{i=1}^L \lambda_i \tilde{\beta}_{t-i} \right) - \frac{2}{L} \sum_{i=1}^L E \tilde{\beta}_t \tilde{\beta}_{t-i} \\
 &= N \text{Var}(\beta) + \sum_{i=1}^L \sum_{j=1}^L E \tilde{\beta}'_{t-i} \lambda'_i \lambda_j \tilde{\beta}_{t-j} - \frac{2}{L} \sum_{i=1}^L E \tilde{\beta}_t \tilde{\beta}_{t-i}, \tag{B3}
 \end{aligned}$$

where we used that $E(\lambda_i) = I/L$.

Next use that

$$\begin{aligned}
 \frac{I}{L} &= E(\lambda'_i) = \sum_{j=1}^L E(\lambda'_i \lambda_j) \\
 &= \sum_{j \neq i} E(\lambda'_i \lambda_j) + E \lambda' \lambda \\
 &= (L-1) E \lambda'_i \lambda_j + E \lambda' \lambda. \tag{B4}
 \end{aligned}$$

It follows that for $i \neq j$,

$$E(\lambda'_i \lambda_j) = \frac{1}{L(L-1)} I - \frac{1}{L-1} E \lambda' \lambda. \tag{B5}$$

The matrix $E \lambda' \lambda$ has zero off-diagonal elements. By symmetry, all on-diagonal elements are the same and are equal to $\sum_{j=1}^L E \lambda(i, j)^2$ for any i , where $\lambda(i, j)$ is element (i, j) of λ . We refer to these diagonal elements as z .

These results give

$$\begin{aligned}
 & \sum_{i=1}^L \sum_{j=1}^L E \tilde{\beta}'_{t-i} \lambda'_i \lambda_j \tilde{\beta}_{t-j} \\
 &= NLz \text{Var}(\beta) + N \left[\frac{1}{L(L-1)} - \frac{1}{L-1} z \right] \text{Var}(\beta) \sum_{i=1}^L \sum_{j \neq i} (\rho_\beta)^{|j-i|} \\
 &= NLz \text{Var}(\beta) + 2N \left[\frac{1}{L(L-1)} - \frac{1}{L-1} z \right] \text{Var}(\beta) \sum_{i=1}^{L-1} (L-i) (\rho_\beta)^i. \tag{B6}
 \end{aligned}$$

We then have

$$\begin{aligned}
 & E\left(\beta_t - \sum_{i=1}^L \lambda_i \beta_{t-i}\right)' \left(\beta_t - \sum_{i=1}^L \lambda_i \beta_{t-i}\right) \\
 &= N \text{Var}(\beta) + NLz \text{Var}(\beta) + 2N \left[\frac{1}{L(L-1)} - \frac{1}{L-1} z \right] \\
 &\quad \times \text{Var}(\beta) \sum_{i=1}^{L-1} (L-i)(\rho_\beta)^i - \frac{2N}{L} \sum_{i=1}^L (\rho_\beta)^i \text{Var}(\beta) \\
 &= N \text{Var}(\beta) + Nz \text{Var}(\beta) \left[L - \frac{2}{L-1} \sum_{i=1}^{L-1} (L-i)(\rho_\beta)^i \right] \tag{B7}
 \end{aligned}$$

$$- \frac{2N}{L(L-1)} \sum_{i=2}^L (i-1)(\rho_\beta)^i \text{Var}(\beta)$$

$$= N \text{Var}(\beta)h, \tag{B8}$$

where

$$\begin{aligned}
 h &= z \left[L - \frac{2}{L-1} \sum_{i=1}^{L-1} (L-i)(\rho_\beta)^i \right] - \frac{2}{L(L-1)} \sum_{i=2}^L (i-1)(\rho_\beta)^i + 1 \\
 &= z \frac{2}{L-1} \sum_{i=1}^{L-1} (L-i)[1 - (\rho_\beta)^i] + \frac{2}{L(L-1)} \sum_{i=2}^L (i-1)[1 - (\rho_\beta)^i]. \tag{B9}
 \end{aligned}$$

Here we used that

$$\sum_{i=1}^{L-1} (L-i) = \sum_{i=2}^L (i-1) = 0.5L(L-1).$$

We then have

$$\text{MSE}^{\text{MODEL}} = N \text{Var}(f) \text{Var}(\beta)h + \text{Var}(u) + \text{Var}(u)E(f'_t \psi_t f_t). \tag{B10}$$

This implies

$$\begin{aligned}
 \text{MSE}^{\text{MODEL}} - \text{MSE}^{\text{RW}} &= N \text{Var}(f)[- \beta^2 - \text{Var}(\beta) + \text{Var}(\beta)h] \\
 &\quad + \text{Var}(f)E(f'_t \psi_t f_t). \tag{B11}
 \end{aligned}$$

This implies

$$\begin{aligned}
 \frac{\text{MSE}^{\text{MODEL}}}{\text{MSE}^{\text{RW}}} &= 1 + \frac{\text{MSE}^{\text{MODEL}} - \text{MSE}^{\text{RW}}}{\text{MSE}^{\text{RW}}} \\
 &= 1 + \frac{-N \text{Var}(f)\beta^2 + \text{Var}(u)E(f'_t \psi_t f_t)}{N \text{Var}(f)[\beta^2 + \text{Var}(\beta)] + \text{Var}(u)} \\
 &\quad + \frac{N \text{Var}(f) \text{Var}(\beta)(h-1)}{N \text{Var}(f)[\beta^2 + \text{Var}(\beta)] + \text{Var}(u)}. \tag{B12}
 \end{aligned}$$

Appendix C

Algebra Section VI

The in-sample fit is then defined as

$$R^2 = 1 - \frac{\text{MSE}^{\text{INSAMPLE}}}{\text{MSE}^{\text{RW}}}, \tag{C1}$$

where

$$\begin{aligned}
 \text{MSE}^{\text{INSAMPLE}} &= \frac{1}{L} \sum_{j=1}^L E(\Delta s_{t-j} - f'_{t-j} \hat{\delta}_t)^2 \\
 &= \frac{1}{L} \sum_{j=1}^L E \left[f'_{t-j} \left(\beta_{t-j} - \sum_{i=1}^L \lambda_i \beta_{t-i} \right) \right. \\
 &\quad \left. + u_{t-j} - f'_{t-j} \psi_t \sum_{i=1}^L u_{t-i} f_{t-i} \right]^2. \tag{C2}
 \end{aligned}$$

This gives

$$\begin{aligned}
 \text{MSE}^{\text{INSAMPLE}} &= \frac{1}{L} \text{Var}(f) \sum_{j=1}^L E \left(\beta_{t-j} - \sum_{i=1}^L \lambda_i \beta_{t-i} \right)' \\
 &\quad \times \left(\beta_{t-j} - \sum_{i=1}^L \lambda_i \beta_{t-i} \right) + \frac{1}{L} \sum_{j=1}^L E \left(u_{t-j} - f'_{t-j} \psi_t \sum_{i=1}^L u_{t-i} f_{t-i} \right)' \\
 &\quad \times \left(u_{t-j} - f'_{t-j} \psi_t \sum_{i=1}^L u_{t-i} f_{t-i} \right). \tag{C3}
 \end{aligned}$$

We have

$$\frac{1}{L} \sum_{j=1}^L E \left(u_{t-j} - f'_{t-j} \psi_t \sum_{i=1}^L u_{t-ij} f_{t-i} \right)' \left(u_{t-j} - f'_{t-j} \psi_t \sum_{i=1}^L u_{t-ij} f_{t-i} \right) \quad (\text{C4})$$

$$\begin{aligned} &= \text{Var}(u) - \frac{2}{L} \sum_{j=1}^L \sum_{i=1}^L E(u_{t-i} u_{t-j} f'_{t-j} \psi_t f_{t-i}) \\ &\quad + \frac{1}{L} E \left[\left(\sum_{i=1}^L u_{t-ij} f_{t-i} \right) \left(\sum_{j=1}^L \psi_t f'_{t-j} f'_{t-j} \psi_t \right) \left(\sum_{i=1}^L u_{t-ij} f_{t-i} \right) \right]. \quad (\text{C5}) \end{aligned}$$

Notice that the second term is nonzero only for $i = j$. Also use that

$$\sum_{i=1}^L f'_{t-i} \psi_t f_{t-i} = \text{diag} \left(\sum_{i=1}^L \psi_t f'_{t-i} f_{t-i} \right) = \text{diag}(I_N) = N. \quad (\text{C6})$$

We then have

$$\frac{1}{L} \sum_{j=1}^L E \left(u_{t-j} - f'_{t-j} \psi_t \sum_{i=1}^L u_{t-ij} f_{t-i} \right)' \left(u_{t-j} - f'_{t-j} \psi_t \sum_{i=1}^L u_{t-ij} f_{t-i} \right) \quad (\text{C7})$$

$$\begin{aligned} &= \text{Var}(u) - \frac{2N}{L} \text{Var}(u) + \frac{1}{L} \text{Var}(u) E \left(\sum_{i=1}^L f'_{t-i} \psi_t f_{t-i} \right) \\ &= \text{Var}(u) - \frac{2N}{L} \text{Var}(u) + \frac{N}{L} \text{Var}(u) = \text{Var}(u) \left(1 - \frac{N}{L} \right). \quad (\text{C8}) \end{aligned}$$

Therefore,

$$\begin{aligned} \text{MSE}^{\text{INSAMPLE}} &= \frac{1}{L} \text{Var}(f) \sum_{j=1}^L E \left(\beta_{t-j} - \sum_{i=1}^L \lambda_i \beta_{t-i} \right)' \\ &\quad \times \left(\beta_{t-j} - \sum_{i=1}^L \lambda_i \beta_{t-i} \right) + \text{Var}(u) \left(1 - \frac{N}{L} \right). \quad (\text{C9}) \end{aligned}$$

We have

$$\begin{aligned}
& \frac{1}{L} \sum_{j=1}^L E \left(\beta_{t-j} - \sum_{i=1}^L \lambda_i \beta_{t-i} \right)' \left(\beta_{t-j} - \sum_{i=1}^L \lambda_i \beta_{t-i} \right) \\
&= N \text{Var}(\beta) - \frac{2}{L^2} \sum_{i=1}^L \sum_{j=1}^L E \tilde{\beta}'_{t-j} \tilde{\beta}_{t-i} + \sum_{i=1}^L \sum_{j=1}^L E \tilde{\beta}'_{t-i} \lambda_i \lambda_j \tilde{\beta}_{t-j} \\
&= N \text{Var}(\beta) - \frac{2N}{L} \text{Var}(\beta) - \frac{2N}{L^2} \text{SUM} \text{Var}(\beta) + N L z \text{Var}(\beta) \\
&+ N \left[\frac{1}{L(L-1)} - \frac{1}{L-1} z \right] \text{Var}(\beta) \text{SUM}, \tag{C10}
\end{aligned}$$

where we used (B6) and

$$\text{SUM} = \sum_{i=1}^L \sum_{j \neq i}^L (\rho_\beta)^{|i-j|} = 2 \sum_{i=1}^{L-1} (L-i) (\rho_\beta)^i. \tag{C11}$$

Note that we can write SUM as

$$\begin{aligned}
\text{SUM} &= -2 \sum_{i=1}^{L-1} (L-i) [1 - (\rho_\beta)^i] + 2 \sum_{i=1}^{L-1} (L-i) \\
&= -2 \sum_{i=1}^{L-1} (L-i) [1 - (\rho_\beta)^i] + L(L-1). \tag{C12}
\end{aligned}$$

Substituting this, we get

$$\frac{1}{L} \sum_{j=1}^L E \left(\beta_{t-j} - \sum_{i=1}^L \lambda_i \beta_{t-i} \right)' \left(\beta_{t-j} - \sum_{i=1}^L \lambda_i \beta_{t-i} \right) = N \text{Var}(\beta) \omega, \tag{C13}$$

where

$$\omega = 2 \left[\frac{L-2}{L^2(L-1)} + \frac{1}{L-1} z \right] \sum_{i=1}^{L-1} (L-i) [1 - (\rho_\beta)^i]. \tag{C14}$$

We then have

$$\text{MSE}^{\text{INSAMPLE}} = N \text{Var}(f) \text{Var}(\beta)\omega + \text{Var}(u) \left(1 - \frac{N}{L}\right). \quad (\text{C15})$$

Therefore,

$$\begin{aligned} \text{MSE}^{\text{INSAMPLE}} - \text{MSE}^{\text{RW}} &= N \text{Var}(f) \\ &\times [-\beta^2 - \text{Var}(\beta) + \text{Var}(\beta)\omega] - \frac{N}{L} \text{Var}(u). \end{aligned} \quad (\text{C16})$$

This implies

$$\begin{aligned} R^2 &= 1 - \frac{\text{MSE}^{\text{INSAMPLE}}}{\text{MSE}^{\text{RW}}} = - \frac{\text{MSE}^{\text{MODEL}} - \text{MSE}^{\text{RW}}}{\text{MSE}^{\text{RW}}} \\ &= \frac{N \text{Var}(f)\beta^2 + (N/L) \text{Var}(u)}{N \text{Var}(f)[\beta^2 + \text{Var}(\beta)] + \text{Var}(u)} \\ &\quad + \frac{N \text{Var}(f) \text{Var}(\beta)(1 - \omega)}{N \text{Var}(f)[\beta^2 + \text{Var}(\beta)] + \text{Var}(u)}. \end{aligned} \quad (\text{C17})$$

Endnotes

Van Wincoop gratefully acknowledges financial support from the Bankard Fund for Political Economy and the National Science Foundation (grant SES-0649442). Bacchetta acknowledges financial support by the National Centre of Competence in Research "Financial Valuation and Risk Management" (NCCR FINRISK) and by the Swiss Finance Institute.

1. See Wolff (1987), Meese and Rogoff (1988), Schinasi and Swamy (1989), and Rossi (2006). In a recent study, Sarno and Valente (2009) show that to achieve the best exchange rate forecast, one needs to continuously change the set of variables used. More generally, there are numerous studies that document structural breaks and regime switching in nominal or real exchange rates.

2. In a closely related paper, Bacchetta and van Wincoop (2009) endogenously derive large time variation in the relationship between exchange rates and fundamentals as a result of incomplete information about very slow moving structural parameters of the economy. In this paper we take the unstable relationship between exchange rates and fundamentals as exogenously given.

3. The variables considered are consistently available over the full sample for the six countries considered. The use of lagged interest rates is justified in Molodtsova and Papell (2009), which evaluates the predictive content of Taylor rule fundamentals for exchange rates. Notice that the vast empirical literature on exchange rates has shown that the precise set of variables is not crucial for the results.

4. Some authors show a better forecasting performance when equations are estimated simultaneously. See, e.g., Mark and Sul (2001), Groen (2005), Cerra and Saxena (2008), Rogoff and Stavrakeva (2008), and Carriero, Kapetanios, and Marcellino (2009).

5. This should therefore not be considered as a true forecasting exercise since the actual future fundamentals are used.

6. There are obviously various ways to evaluate forecasting performance. In this paper, we restrict ourselves to the MSE ratio. Many recent papers consider tests taking into account small-sample biases, e.g., following Clark and West (2006). It is not necessary to consider such an adjustment in our context since our objective is to compare actual data to the data generated by a model with the same sample size.

7. We have scaled the original fundamentals by a constant number so that the estimated coefficient is equal to one in a regression over the whole sample. The next section explains this normalization in more detail.

8. Chen and Rogoff (2003) report that world commodity prices are an important determinant for the real exchange rate of several major commodity exporters, such as Australia, Canada, and New Zealand. Chen (2004) finds that augmenting standard monetary models for the nominal exchange rate of commodity currencies with commodity export prices improves the in-sample fit. In terms of out-of-sample fit, the improvement is mixed and depends on the specification chosen. Chen, Rogoff, and Rossi (2008) show that the reverse relationship is much stronger: exchange rates of commodity currencies contain useful information to predict global commodity prices.

9. Australia and New Zealand have experienced shorter floating exchange rate episodes than the other currencies. In order to be able to compare the MSE ratio from the three currencies, we choose the largest possible common sample, which starts in January 1986 and ends in December 2008.

10. The reason is that observed fundamentals have limited explanatory power as measured by the low R^2 for long data samples. Therefore, the standard deviation of u_t is quite close to the standard deviation of Δs_t .

11. The unconditional variance, which is $\sigma_\beta^2/(1 - \rho_\beta^2)$, is then one as well. For a given ρ_β , we then set $\sigma_\beta^2 = 1 - \rho_\beta^2$.

12. That small samples are the reason for the failure of fundamentals-based models relative to the random walk model is often mentioned in the literature (e.g., Engel and West 2005).

References

- Alquist, Ron, and Menzie D. Chinn. 2008. "Conventional and Unconventional Approaches to Exchange Rate Modelling and Assessment." *International Journal of Finance and Economics* 13:2–13.
- Bacchetta, Philippe, and Eric van Wincoop. 2009. "On the Unstable Relationship between Exchange Rates and Macroeconomic Fundamentals." Working Paper no. 15008, NBER, Cambridge, MA.
- Carriero, Andrea, George Kapetanios, and Massimiliano Marcellino. 2009. "Forecasting Exchange Rates with a Large Bayesian VAR." *International Journal of Forecasting* 25:400–417.
- Cerra, Valerie, and Sweta C. Saxena. 2008. "The Monetary Model Strikes Back: Evidence from the World." Working Paper no. 08/73, International Monetary Fund, Washington, DC.
- Chen, Yu-chin. 2004. "Exchange Rates and Fundamentals: Evidence from Commodity Economies." Manuscript, University of Washington.
- Chen, Yu-chin, and Kenneth Rogoff. 2003. "Commodity Currencies." *Journal of International Economics* 60:133–60.
- Chen, Yu-chin, Kenneth Rogoff, and Barbara Rossi. 2008. "Can Exchange Rates Forecast Commodity Prices?" Working Paper no. 13901, NBER, Cambridge, MA.
- Cheung, Yin-Wong, Menzie D. Chinn, and Antonio Garcia Pascual. 2005. "Empirical Exchange Rate Models of the Nineties: Are Any Fit to Survive?" *Journal of International Money and Finance* 24, no. 7:1150–75.
- Clark, Todd E., and Kenneth D. West. 2006. "Using Out-of-Sample Mean Square Prediction Errors to Test the Martingale Difference Hypothesis." *Journal of Econometrics* 135:155–86.

- Engel, Charles, and Kenneth D. West. 2005. "Exchange Rates and Fundamentals." *Journal of Political Economy* 113:485–517.
- Groen, Jan J. J. 2005. "Exchange Rate Predictability and Monetary Fundamentals in a Small Multi-country Panel." *Journal of Money, Credit and Banking* 37, no. 3:495–516.
- Mark, Nelson C., and Donggyu Sul. 2001. "Nominal Exchange Rates and Monetary Fundamentals: Evidence from a Small Post-Bretton Woods Panel." *Journal of International Economics* 53:29–52.
- Meese, Richard A., and Kenneth Rogoff. 1983a. "Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?" *Journal of International Economics* 14:345–73.
- . 1983b. "The Out-of-Sample Failure of Empirical Exchange Rate Models: Sampling Error or Misspecification?" In *Exchange Rates and International Macroeconomics*, ed. J. Frenkel, 67–105. Chicago: University of Chicago Press.
- . 1988. "Was It Real? The Exchange Rate–Interest Differential Relation over the Modern Floating-Rate Period." *Journal of Finance* 43:933–48.
- Molodtsova, Tanya, and David H. Papell. 2009. "Out-of-Sample Exchange Rate Predictability with Taylor Rule Fundamentals." *Journal of International Economics* 77:167–80.
- Rogoff, Kenneth, and Vania Stavrakeva. 2008. "The Continuing Puzzle of Short Horizon Exchange Rate Forecasting." Working Paper no. 14071, NBER, Cambridge, MA.
- Rossi, Barbara. 2006. "Are Exchange Rates Really Random Walks? Some Evidence Robust to Parameter Instability." *Macroeconomic Dynamics* 10:20–38.
- Sarno, Lucio, and Giorgio Valente. 2009. "Exchange Rates and Fundamentals: Footloose or Evolving Relationship?" *Journal of the European Economic Association* 7:786–830.
- Schinasi, Garry J., and P. A. V. B. Swamy. 1989. "The Out-of-Sample Forecasting Performance of Exchange Rate Models When Coefficients Are Allowed to Change." *Journal of International Money and Finance* 8:375–90.
- Wolff, Christian C. P. 1987. "Time-Varying Parameters and the Out-of-Sample Forecasting Performance of Structural Exchange Rate Models." *Journal of Business and Economic Statistics* 5:87–97.