

NATIONAL BANK OF POLAND WORKING PAPER №. 123

Real exchange rate forecasting: a calibrated
half-life PPP model can beat the random walk

Michele Ca' Zorzi, Michał Rubaszek

Michele Ca' Zorzi - European Central Bank.

E-mail: michele.cazorzi@ecb.europa.eu

Michał Rubaszek - National Bank of Poland, Warsaw School of Economics.

E-mail: michal.rubaszek@nbp.pl.

The views expressed in this paper are those of the authors and not necessarily those of the National Bank of Poland or the European Central Bank.

Design:

Oliwka s.c.

Layout and print:

NBP Printshop

Published by:

National Bank of Poland

Education and Publishing Department

00-919 Warszawa, 11/21 Świętokrzyska Street

phone: +48 22 653 23 35, fax +48 22 653 13 21

© Copyright by the National Bank of Poland, 2012

ISSN 2084-624X

<http://www.nbp.pl>

Contents

Abstract	3
1 Introduction	5
2 The models	7
3 Empirical evidence	8
4 Analytical interpretation of the results	10
5 Conclusions	13

List of Tables

1	Mean Squared Forecast Errors (MSFEs)	16
2	Correlation of forecast and realized changes of real exchange rates . .	16

List of Figures

1	Real exchange rates (2010 = 100)	17
2	Mean Squared Forecast Errors	18
3	Theoretical Mean Squared Forecast Errors	19

Abstract

This paper brings two new insights into the Purchasing Power Parity (PPP) debate. First, even if PPP is thought to hold only in the long run, we show that a half-life PPP model outperforms the random walk in real exchange rate forecasting, also at short-term horizons. Second, we show that this result holds as long as the speed of adjustment to the sample mean is imposed and not estimated. The reason is that the estimation error of the pace of convergence distorts the results in favor of the random walk model, even if the PPP holds in the long-run.

Keywords: Exchange rate forecasting; purchasing power parity; half-life.

JEL classification: C32; F31; F37

Non-technical summary

There is broad agreement between policy makers and academics that beating the random walk in forecasting nominal exchange rates is very difficult, if at all possible. Views are instead less unanimous on whether real exchange rates can be forecast.

In this paper we suggest that forecasting changes in real exchange rates is equally important if one wants to form a view on the price competitiveness and export prospects of a given country. The task that we set ourselves is to investigate whether forecasting real effective exchange rate of major world currencies is achieved more accurately with an economic-theory-based model or with a naive random walk. The standard theoretical reference on this issue is the PPP hypothesis, which suggests that the relative price of two identical, domestic and foreign, baskets of goods is constant when expressed in a common currency. Although PPP is one of the most prominent theories in economics, it remains highly controversial, as it is thought to fail in the short-run. As for the long run it is generally recognized that mean reversion to the PPP implied rate is a factor at play.

In this paper we add to the debate on PPP by presenting new evidence in support of PPP for nine major world currencies. We show that a trivially simple model, which assumes slow convergence of the real exchange rate to PPP implied level, generally outperforms, and in many cases in a significant way, the random walk model in terms of real exchange rate forecasting. What is remarkable in comparison to earlier studies, we show that this is true also for short-term horizons.

The article also presents an additional insight which is very telling. We indicate that when the (slow-adjusting) PPP model is estimated, it performs poorly: the outcome is generally worse than the random walk model. There is, nonetheless, a simple solution that is extremely easy to apply in practice. We recommend imposing the speed of mean reversion to PPP at values consistent with the duration of the half-life in the range between 3 and 5 years (in line with the literature). We show that if one does that the real exchange rate forecasts turn out to be much better than those derived with the random walk model. The rest of the paper is devoted to understanding this result from a theoretical point of view. We show that for persistent processes the forecast error attributed to estimation is likely to overwhelm the one caused by the mis-specification of imposing a calibrated model.

Overall we find that the analysis is encouraging on the usefulness of exchange rates theory: by choosing a different battlefield, i.e. real exchange rate forecasting, a theory based model outperforms the random walk.

1 Introduction

Following the seminal papers by Meese and Rogoff (1983a,b), a consensus view has emerged that economic theory is of very little help in forecasting exchange rates. Although in the mid-1990s Mark (1995) and Chinn and Meese (1995) claimed that the random walk (RW) model could be at least beaten at longer horizons, this more optimistic perspective was short-lived and vigorously contested (see in particular Cheung et al., 2005). Looking forward, as stated by (Rogoff, 2009) the unpredictability of exchange rates is likely to remain the consensus view for the conceivable future. The aim of this paper is not to challenge this assessment but to signal some promising avenues of research. While sharing the fascination and desire to understand the underlying forces of exchange rates, in what follows we propose that researchers could deal first with an apparently less ambitious but still key assignment: i.e. real exchange rate (RER) forecasting. Not only this may be easier, it could also be more relevant from a macroeconomic perspective, if one believes, as we do, that to assess a country's outlook, the relevant concept is price competitiveness and not the level of the exchange rate. The obsession with nominal exchange rates probably explains why only a handful of studies have investigated the predictability of RERs. The exceptions are the papers by Meese and Rogoff (1988), Mark and Choi (1997) and Pavlidis et al. (2011) although they reach opposite conclusions: the first rather skeptical and the two last more positive on the scope for RER forecasting.

The standard theoretical reference on RER is the PPP hypothesis, one of the most prominent and controversial theories in the history of economic thinking. In their review of the PPP debate Taylor and Taylor discuss how the consensus has shifted for and against PPP over time: in their assessment the common view is now back to what had prevailed before the 1970s, i.e. "that short run PPP does not hold, that long-run PPP may hold in the sense that there is significant mean reversion of the real exchange rate, although there may be factors impinging on the equilibrium RER through time" (Taylor and Taylor, 2004, p. 154). The empirical literature that conducts unit root tests to evaluate the mean-reversion of RERs usually finds that it is not possible to reject the null of RERs non-stationarity. The evidence, however, is not conclusive owing to the low power of the tests for persistent processes. The "PPP puzzle" literature, which estimates the speed of mean-reversion of RERs, generally concludes that it takes between 3 and 5 years to halve PPP deviations (see Rogoff, 1996; Kilian and Zha, 2002).

This paper adds to the above literature in two aspects. First, we show that a calibrated half-life PPP, which postulates a gradual adjustment of RER to the PPP level, outperforms the RW model in forecasting real exchange rates, also in the

short-term horizon. We claim that this finding provides new evidence in the PPP debate. Second, we show that the calibrated model also outperforms its estimated counterpart. The reason is that in the case of persistent process the estimation error is so high.

The rest of the paper is structured as follows. Section 2 outlines the alternative models that we shall use in our exchange rate forecasting competition. In section 3 we provide some empirical support for the PPP hypothesis using monthly data for real effective exchange rates of major currencies for the period between January 1975 and March 2012. In particular, we show that a calibrated model, which assumes slow convergence to PPP, strongly outperforms the RW at both long and short-term horizons. Finally in Section 4 we provide an analytical investigation of our empirical findings, pointing to the important role of the estimation error.

2 The models

Let us define the log of the real exchange rate as $y_t \equiv s_t + p_t - p_t^*$, where s_t is the log of the nominal exchange rate expressed as the foreign currency price of a unit of domestic currency, and p_t and p_t^* are the logs of home and foreign price levels, respectively. Let us also assume that the DGP for y_t is a simple autoregression (AR) of the form:

$$(y_t - \mu) = \rho(y_{t-1} - \mu) + \epsilon_t, \epsilon_t \sim \mathcal{N}(0, \sigma^2) \quad (1)$$

with $|\rho| < 1$ measuring the speed of reversion to μ , which is interpreted as the level of PPP. As mentioned in the introduction, the consensus view is that the half-life of deviations from the PPP:

$$hl = \log(0.5)/\log(\rho). \quad (2)$$

is somewhere between 3 and 5 years. This view implicitly assumes that RERs are mean reverting and hence predictable. We show that this claim is generally justified by comparing the accuracy of RER forecasts derived from the following competing models, which are a specific form of (1).

The first model is a random walk, for which the h step ahead forecast is:

$$y_{T+h|T}^{RW} = y_T. \quad (3)$$

The next two models assume that the half-life amounts to 3 or 5 years (HL3 and HL5), thus RERs converge to their sample mean values at pace $\bar{\rho}$ consistent with the duration of the half-life in line with (2). The h step ahead forecast is:

$$y_{T+h|T}^{HL} = \bar{\mu} + \bar{\rho}^h(y_T - \bar{\mu}), \quad (4)$$

where $\bar{\mu}$ is the sample mean from the last R observations. The last competitor is the AR model of the form (1) for which:

$$y_{T+h|T}^{AR} = \hat{\mu} + \hat{\rho}^h(y_T - \hat{\mu}), \quad (5)$$

where $\hat{\mu}$ and $\hat{\rho}$ are OLS estimates on the basis of the last R observations.

3 Empirical evidence

To assess the predictability of RERs we gathered monthly data for nine major currencies of the following countries: Australia (AUD), Canada (CAD), euro area (EUR), Japan (JPY), Mexico (MXN), New Zealand (NZD), Switzerland (CHF), the United Kingdom (GBP) and the United States (USD) for the period between 1975:1 and 2012:3. For all currencies we model (narrow) real effective exchange rates as calculated by the Bank for International Settlements (Klau and Fung, 2006). The values of the analyzed series are presented in Figure 1.

The out-of-sample forecast performance is analyzed for horizons ranging from one up to sixty months ahead, whereas the evaluation is based on data from the period 1990:1 to 2012:3. The models are estimated using rolling samples of 15 years ($R = 180$ months). The first set of forecasts is elaborated with the rolling sample 1975:1-1989:12 for the period 1990:1-1994:12. This procedure is repeated with the rolling samples ending in each month from the period 1990:2-2012:2. Since the data available end in 2012:3, the 1-month ahead forecasts are evaluated on the basis of 267 observations, 2-month ahead forecasts on the basis of 266 observations, and 60-month ahead forecasts on the basis of 208 observations.

The forecasting performance is measured with two standard statistics: the mean squared forecast errors (MSFEs) and the correlation coefficient between forecast and realized RER changes. Table 1 and Figure 2 present the values of MSFEs. As is generally done in the forecasting literature, we report the actual MSFEs values for the RW model, while for the remaining models the numbers are expressed as ratios, so that values below unity indicate that a given model dominates the RW. We also test the null of equal forecast accuracy with the two-sided Diebold and Mariano (1995) test.

In terms of the MFSE criterion the two HL model-based forecasts are significantly better than the RW for seven out of nine currencies (EUR, MXN, NZD, CHF, GBP, USD, JPY). In the specific case of HL5, we find that the MSFEs are on average 9% and 23% lower than that from the RW model at the two and five-year horizon, respectively. Looking at the specific currencies, both H3 and H5 model-based forecasts are much more precise than those based on the AR model for the following five currencies (CAD, EUR, JPY, GBP and USD) while the outcomes are broadly comparable for the other four currencies. Particularly interesting is that HL models are able to outperform other models also at short-term horizons - at which the RW model is generally thought to be very difficult to beat. The short-term forecasts from the HL models outperform significantly those based on the RW model in the case of EUR, MXN and NZD, at the same time being broadly comparable for the other currencies. At one year horizon, the MSFEs from the HL5 model are on

average by 3% and 12% lower than those from the RW and AR models, respectively. Finally, at short-term horizons the AR model performs particularly poorly compared to both the HL and RW models.

Further evidence that the HL models beat the alternatives can be found using our second criterion, which consists in computing the correlation coefficient between the realized and forecast changes of RERs:

$$r_{M,h} = \text{cor}(y_{T+h|T}^M - y_T, y_{T+h} - y_T), \quad (6)$$

where M stands for the model name. Note that (3) and (4) imply that $r_{RW,h}$ is zero and $r_{HL,h}$ does not depend on the duration of the half-life: for that reason in Table 2 we report only the results for two models, a common HL and the AR model. The table shows that the correlation coefficients for the HL model are generally positive for all currencies at all horizons, except for the AUD. The average value of $r_{HL,h}$ also increases with the forecast horizon: from just 0.04 for the one-month ahead forecasts to 0.53 for the five-year ahead forecasts. In the case of the AR model the results are again rather disappointing: MXN is the only currency with a positive $r_{AR,h}$ throughout the forecast horizon. Moreover, the average value of $r_{AR,h}$ is positive only for horizons above two years. Finally, at all horizons $r_{AR,h}$ is visibly lower than $r_{HL,h}$.

To sum up, there appears to be convincing evidence that RERs of major currencies are mean reverting and forecastable. A puzzling question, however, remains: if the true DGP is given by (1), why is the AR model performing so poorly compared to the RW model? The short-answer is that estimation error plays a key role here: even if we estimate the speed of adjustment using fairly long windows (15 years), the error is large enough to distort the results in favor of the RW model even if PPP holds in the long-run.

4 Analytical interpretation of the results

In this section we show analytically the point that we have just stated. Let assume that the DGP for y_t is given by (1) so that the unbiased and efficient forecast is:

$$y_{T+h|T} = \mu + \rho^h(y_T - \mu), \quad (7)$$

and the variance of the forecast error:

$$E\{(y_{T+h} - y_{T+h|T})^2\} = \sigma^2 \frac{1 - \rho^{2h}}{1 - \rho^2}. \quad (8)$$

Here, the only source of forecast errors is the existence of the unobserved random term, whose future values are unknown (random error). In the case of forecasts given by equations (3)-(5) the variance of forecast errors is higher than that in (8) because the coefficients μ and ρ are unknown and have to be estimated (estimation error) or calibrated (calibration error).

Let us decompose the variance of the forecast error from a generic model $M \in \{RW, HL, AR\}$ into three components:

$$\begin{aligned} E\{(y_{T+h} - y_{T+h|T}^M)^2\} &= E\{(y_{T+h} - y_{T+h|T})^2\} + E\{(y_{T+h|T} - y_{T+h|T}^M)^2\} + \\ &+ 2E\{(y_{T+h} - y_{T+h|T})(y_{T+h|T} - y_{T+h|T}^M)\}, \end{aligned} \quad (9)$$

The value of the first component, which is given by (8), represents the random error. The second component quantifies the mis-specification, estimation or calibration error. Finally, if we cannot forecast future shocks the third component is null and can be disregarded.

In what follows we provide the analytical expressions for the second component, which is what matters to assess the relative performance of our competing models.

In the case of the RW model the error equals:

$$y_{T+h|T} - y_{T+h|T}^{RW} = (1 - \rho^h)(y_T - \mu), \quad (10)$$

and thus:

$$E\{(y_{T+h|T} - y_{T+h|T}^{RW})^2\} = (1 - \rho^h)^2 \times E\{(y_T - \mu)^2\}, \quad (11)$$

where:

$$E\{(y_T - \mu)^2\} = \frac{\sigma^2}{1 - \rho^2}$$

For the HL model, the combination of (4) and (7) yields:

$$y_{T+h|T} - y_{T+h|T}^{HL} = (\rho^h - \bar{\rho}^h)(y_T - \mu) - (1 - \bar{\rho}^h)(\bar{\mu} - \mu). \quad (12)$$

The first term describes the forecast error caused by the wrong calibration of the parameter ρ , which determines the adjustment speed to PPP; the second term is instead the forecast error related to the estimation of μ , i.e. the constant, which approximates the PPP value on the basis of the sample of size R . The resulting variance is:

$$\begin{aligned} E\{(y_{T+h|T} - y_{T+h|T}^{HL})^2\} &= (\rho^h - \bar{\rho}^h)^2 \times E\{(y_T - \mu)^2\} + (1 - \bar{\rho}^h)^2 \times \\ &\times E\{(\bar{\mu} - \mu)^2\} - 2(\rho^h - \bar{\rho}^h)(1 - \bar{\rho}^h) \times E\{(y_T - \mu)(\bar{\mu} - \mu)\} \end{aligned} \quad (13)$$

where:

$$\begin{aligned} E\{(\bar{\mu} - \mu)^2\} &= \frac{\sigma^2}{1 - \rho^2} \times \frac{1}{R^2} \times (R + 2 \sum_{j=1}^{R-1} (R - j)\rho^j) \\ E\{(y_T - \mu)(\bar{\mu} - \mu)\} &= \frac{\sigma^2}{1 - \rho^2} \times \frac{1}{R} \times \frac{1 - \rho^R}{1 - \rho}. \end{aligned}$$

Finally, as derived in Fuller and Hasza (1980), for the AR model the variance of the estimation error is approximately:

$$E\{(y_{T+h|T} - y_{T+h|T}^{AR})^2\} \simeq \sigma^2 \times \frac{1}{R} \times \left[h^2 \rho^{2(h-1)} + \left(\frac{1 - \rho^h}{1 - \rho} \right)^2 \right]. \quad (14)$$

Given equations (8)-(14), the assumptions for the DGP coefficients (μ , ρ and σ) and the sample size (R), one can calculate the theoretical value of MSFE for all competing models (RW, HL and AR) at different forecast horizons ($h = 1, 2, \dots, H$). It is worth noting that the theoretical MSFEs of all models do not depend on the value of μ and are proportional to the value of σ^2 . Consequently, the relative MSFEs depend solely on the convergence coefficient ρ , the sample size R and forecast horizon h .

In what follows we consider values of ρ corresponding to a half-life varying from one to ten years, the sample size of 180 monthly observations and forecast horizons up to 60 months ahead. These values correspond to the empirical analysis described in Section 3. The results are presented in Figure 3, where the theoretical MSFEs values of a given model are shown as a ratio of the MSFEs of the RW model.

The analytical results, which were cross-checked with Monte Carlo simulations, suggest that the HL3 and HL5 model-based forecasts are considerably more accurate than those from RW models if the half-life of DGP is below 5 years. For more persistent deviations than that, the RW model would outperform other models. The HL3 and HL5 model-based forecasts are instead more accurate than those from the AR model if the DGP half-life is above one year. Finally in terms of point forecast accuracy, the AR model outperforms the RW model only for relatively low values of the DGP half-life, i.e. not exceeding three years. The estimation error of the AR model is especially severe for more persistent processes.

The above results tell the following story: even if the true DGP is a simple autoregression with the duration of the half-life between three and five years, an estimated AR model usually will not outperform the RW model in forecasting. This result is explained by the estimation forecast error of the AR model, which outweighs the accuracy loss of choosing the mis-specified RW model. The remedy is, however, simple: just employ a reasonably calibrated HL model that assumes a gradual mean reversion to the sample mean.

5 Conclusions

In this paper we have shown that a calibrated half-life model is remarkably successful in forecasting RERs at both long and short term horizons. We highlight the important role estimation error may play, as it could lead to wrong conclusions even if one has identified the correct model. Overall, we find the analysis is encouraging on the usefulness of exchange rate theory: by choosing a different battlefield, i.e. RERs forecasting, a simple modification of PPP theory generally outperforms, and by far, the RW model. Looking forward, RERs forecasting appear a relevant and promising avenue of research: it may prove nonetheless challenging to push the analysis beyond the major currencies for which PPP is a natural theoretical framework. Our suggestion to anyone pursuing this research to the case of emerging countries where price convergence plays a role is however clear: beware of the role of the estimation error and consider also whether simple calibrated models are competitive.

References

- Cheung, Y.-W., Chinn, M. D., and Pascual, A. G. (2005). Empirical exchange rate models of the nineties: Are any fit to survive? *Journal of International Money and Finance*, 24(7):1150–1175.
- Chinn, M. D. and Meese, R. A. (1995). Banking on currency forecasts: How predictable is change in money? *Journal of International Economics*, 38(1-2):161–178.
- Diebold, F. X. and Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3):253–63.
- Fuller, W. A. and Hasza, D. P. (1980). Predictors for the first-order autoregressive process. *Journal of Econometrics*, 13(2):139 – 157.
- Kilian, L. and Zha, T. (2002). Quantifying the uncertainty about the half-life of deviations from PPP. *Journal of Applied Econometrics*, 17(2):107–125.
- Klau, M. and Fung, S. S. (2006). The new BIS effective exchange rate indices. *BIS Quarterly Review*.
- Mark, N. C. (1995). Exchange rates and fundamentals: evidence on long-horizon predictability. *American Economic Review*, 85(1):201–18.
- Mark, N. C. and Choi, D.-Y. (1997). Real exchange-rate prediction over long horizons. *Journal of International Economics*, 43(1-2):29–60.
- Meese, R. and Rogoff, K. (1983a). The out-of-sample failure of empirical exchange rate models: Sampling error or misspecification? In Frenkel, J. A., editor, *Exchange Rates and International Macroeconomics*, NBER Chapters, pages 67–112. National Bureau of Economic Research, Inc.
- Meese, R. A. and Rogoff, K. (1983b). Empirical exchange rate models of the seventies : Do they fit out of sample? *Journal of International Economics*, 14(1-2):3–24.
- Meese, R. A. and Rogoff, K. (1988). Was it real? the exchange rate-interest differential relation over the modern floating-rate period. *Journal of Finance*, 43(4):933–48.
- Pavlidis, E. G., Paya, I., and Peel, D. A. (2011). Forecast evaluation of nonlinear models: The case of long-span real exchange rates. *Journal of Forecasting*, <http://dx.doi.org/10.1002/for.1247>.

- Rogoff, K. (1996). The purchasing power parity puzzle. *Journal of Economic Literature*, 34(2):647–668.
- Rogoff, K. (2009). Exchange rates in the modern floating era: what do we really know? *Review of World Economics (Weltwirtschaftliches Archiv)*, 145(1):1–12.
- Taylor, A. M. and Taylor, M. P. (2004). The purchasing power parity debate. *Journal of Economic Perspectives*, 18(4):135–158.

Table 1: Mean Squared Forecast Errors (MSFEs)

h	RW	HL3	HL5	AR1	RW	HL3	HL5	AR1	RW	HL3	HL5	AR1
	AUD				CAD				EUR			
1	0.05	1.03	1.01	1.02*	0.05	1.04	1.02	1.03*	0.02	1.00	1.00	1.04**
6	0.44	1.07	1.03	1.06*	0.46	1.10	1.03	1.11**	0.23	0.96	0.96	1.12*
12	0.82	1.13	1.06	1.09*	0.89	1.17	1.06	1.20**	0.49	0.90	0.92*	1.14*
24	1.53	1.22**	1.10*	1.09	1.68	1.24*	1.10	1.19*	1.03	0.80*	0.83**	1.18**
36	2.10	1.25*	1.12	1.06	2.24	1.18	1.06	1.20*	1.67	0.72**	0.77**	1.13*
60	3.00	1.22	1.11	1.06	3.17	0.99	0.94	1.45**	2.84	0.58**	0.66**	0.91
	JPY				MXN				NZD			
1	0.06	1.01	1.00	1.01	0.07	0.99	0.99	0.99	0.11	1.00	1.00	1.03*
6	0.59	0.99	0.98	1.04	0.61	0.94	0.96*	0.92	0.80	0.96	0.96	1.06
12	1.00	1.00	0.97	1.10*	1.11	0.89*	0.92*	0.87*	1.42	0.90*	0.92*	1.03
24	2.34	0.92	0.91	1.17**	2.73	0.79**	0.84**	0.78*	2.53	0.77**	0.83**	0.89*
36	3.53	0.86	0.86	1.19**	4.19	0.72**	0.78**	0.74**	2.85	0.65**	0.74**	0.74**
60	3.12	0.97	0.89	1.19*	3.72	0.73*	0.74**	0.72*	2.64	0.55**	0.64**	0.62**
	CHF				GBP				USD			
1	0.02	1.00	1.00	1.06	0.02	1.00	1.00	1.02	0.03	1.00	1.00	1.03*
6	0.12	0.98	0.98	1.22*	0.15	0.97	0.97	1.04	0.23	0.96	0.96	1.10*
12	0.25	0.97	0.97	1.16	0.29	0.95	0.95	1.03	0.45	0.94	0.93	1.21**
24	0.50	0.84*	0.88*	0.99	0.49	0.85*	0.87*	0.98	0.92	0.84	0.84*	1.21*
36	0.72	0.73**	0.80**	0.78**	0.56	0.77**	0.82**	0.93	1.21	0.71*	0.72**	1.13
60	0.79	0.65**	0.72**	0.69**	0.55	0.59**	0.67**	0.70**	1.32	0.45**	0.53**	0.91

Notes: For the RW model MSFEs are reported in levels (multiplied by 100), whereas for the remaining methods they appear as the ratios to the corresponding MSFE from the RW model. Asterisks ** and * denote the rejection of the null of the Diebold and Mariano (1995) test, stating that the MSFE from RW are not significantly different from the MSFE of a given model, at 1%, 5% significance level, respectively.

Table 2: Correlation of forecast and realized changes of real exchange rates

h	AUD	CAD	EUR	JPY	MXN	NZD	CHF	GBP	USD	av.
	Half-life models									
1	-0.04	-0.03	0.06	0.04	0.11	0.06	0.08	0.06	0.07	0.04
6	-0.01	0.02	0.21	0.15	0.30	0.23	0.21	0.19	0.22	0.17
12	-0.05	0.01	0.32	0.18	0.42	0.36	0.25	0.26	0.29	0.23
24	-0.10	0.02	0.46	0.31	0.60	0.56	0.50	0.43	0.43	0.36
36	-0.12	0.11	0.55	0.38	0.66	0.73	0.66	0.53	0.56	0.45
60	-0.20	0.26	0.71	0.31	0.61	0.81	0.74	0.71	0.78	0.53
	AR model									
1	-0.04	-0.09	-0.19	-0.04	0.13	-0.07	-0.13	-0.03	-0.09	-0.06
6	-0.07	-0.22	-0.29	-0.10	0.33	-0.05	-0.21	0.02	-0.17	-0.08
12	-0.06	-0.31	-0.24	-0.22	0.44	0.05	-0.04	0.08	-0.24	-0.06
24	0.00	-0.14	-0.22	-0.33	0.60	0.35	0.22	0.19	-0.13	0.06
36	0.07	-0.09	-0.06	-0.32	0.64	0.61	0.52	0.26	0.05	0.19
60	0.10	-0.43	0.32	0.04	0.63	0.78	0.64	0.58	0.34	0.33

Figure 1: Real exchange rates (2010 = 100)


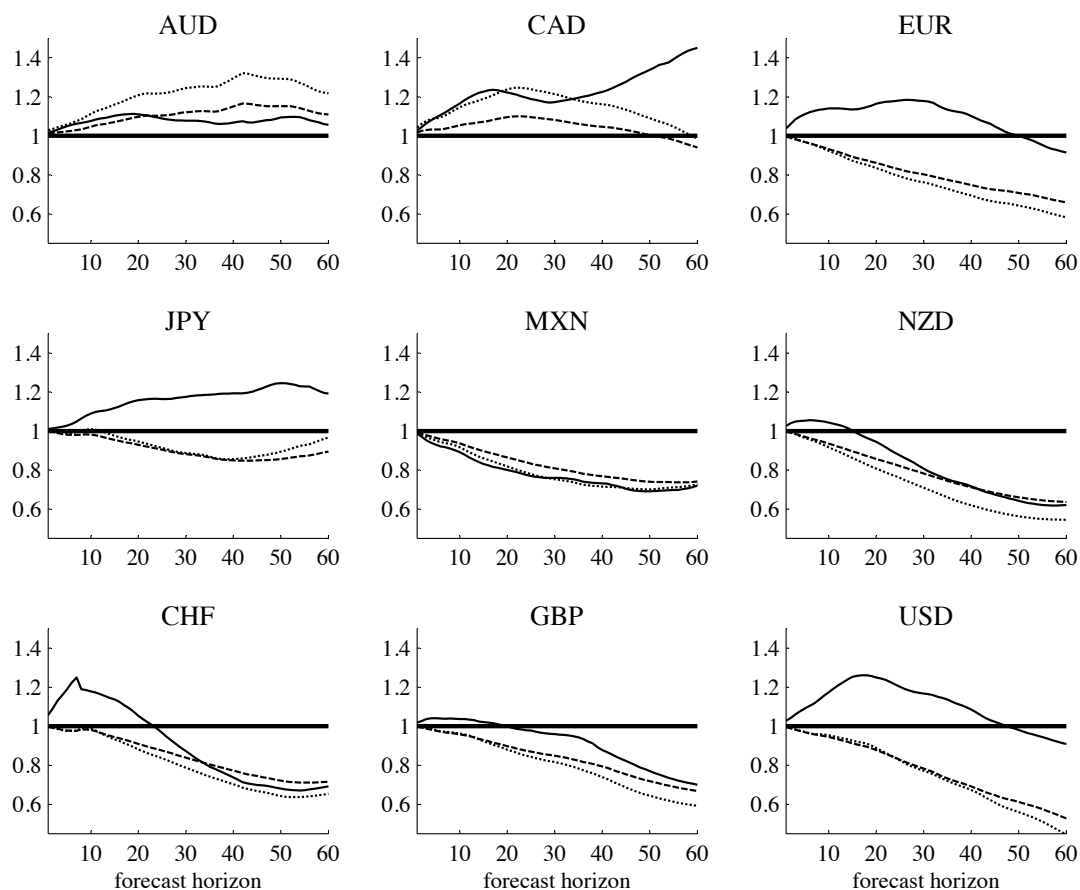
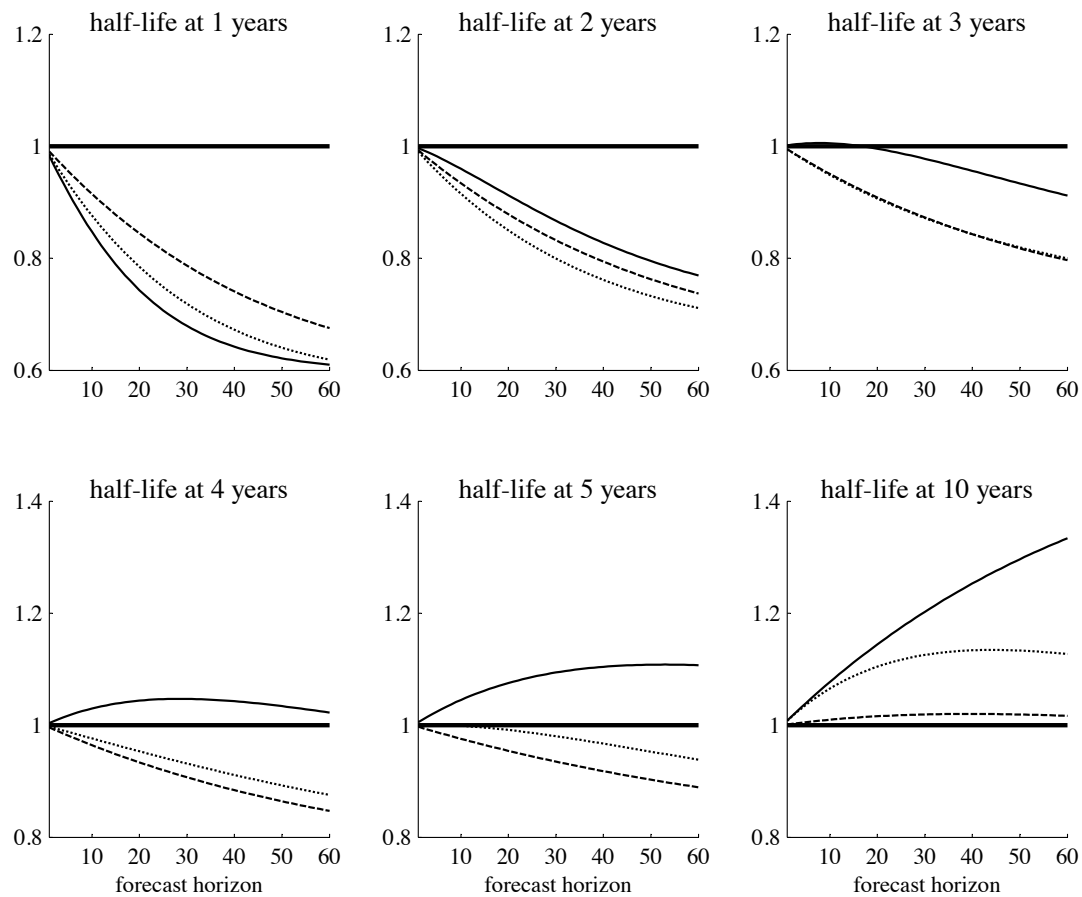


Figure 2: Mean Squared Forecast Errors



Notes: Each line represents the ratio of MSFE from a given method to MSFE from the random walk, where values below unity indicate better accuracy of point forecasts. The straight, dashed and dotted lines stand for AR, HL3 and HL5, respectively. The forecast horizon is expressed in months.

Figure 3: Theoretical Mean Squared Forecast Errors



Notes: Each line represents the ratio of MSFE from a given method to MSFE from the random walk, where values below unity indicate better accuracy of point forecasts. The straight, dashed and dotted lines stand for AR, HL3 and HL5, respectively. The forecast horizon is expressed in months.