

Global Research Unit Working Paper #2018-025

Exchange rate forecasting on a napkin

Michele Ca' Zorzi, ECB
Michal Rubaszek, SGH Warsaw School of
Economics

© 2018 by Ca' Zorzi & Rubaszek. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Exchange rate forecasting on a napkin

Michele Ca' Zorzi
European Central Bank

Michał Rubaszek
SGH Warsaw School of Economics

Abstract

This paper shows that there are two regularities in foreign exchange markets in advanced countries with flexible regimes. First, real exchange rates are mean-reverting, as implied by the Purchasing Power Parity model. Second, the adjustment takes place via nominal exchange rates. These features of the data can be exploited, even on the back of a napkin, to generate nominal exchange rate forecasts that outperform the random walk. The secret is to avoid estimating the pace of mean reversion and assume that relative prices are unchanged. Direct forecasting or panel data techniques are better than the random walk but fail to beat this simple calibrated model.

Keywords: exchange rates; forecasting; Purchasing Power Parity; panel data; mean reversion.

JEL classification: C32; F31; F37; F41.

Non-technical summary

The international finance literature has documented two important regularities in foreign exchange markets. First, there is ample evidence that, for developed countries, real exchange rates are reverting to the level implied by the Purchasing Power Parity (PPP) theory. Second, for flexible currency regimes the adjustment process is mainly driven by the nominal exchange rate. At the same time most of the recent articles remain skeptical that one can outperform the random walk (RW) in nominal exchange rate forecasting.

In this paper we claim that the two above in-sample regularities of foreign exchange markets can be exploited to infer out-of-sample movements of major currency pairs. To prove this thesis we proceed as follows:

1. We begin by presenting robust (in-sample) evidence that, for major currency pairs, long-run PPP holds and that the nominal exchange rate is the main driver of this adjustment process.
2. We then evaluate a battery of models that aim to exploit these in-sample regularities for forecasting purposes. The winner of the forecasting race is a calibrated PPP model, which just assumes that the real exchange rate gradually returns to its sample mean, completing half of the adjustment in 3 years, and that the adjustment is only driven by the nominal exchange rate. This approach is so simple that it can be implemented even on the back of a napkin in two steps. Step 1 consists in calculating the initial real exchange rate misalignment with an eyeball estimate of what is the distance from the sample mean. Step 2 consists in recalling that, according to this model, one tenth of the required adjustment is achieved by the nominal exchange rate in the first 6 months, one fifth in one year, just over a third in two years and exactly half after 3 years.
3. We highlight that severe problems arise when attempting to carry out more sophisticated approaches, such as estimating the pace of mean reversion of the real exchange rate or forecasting relative inflation. Among the estimated approaches, we find that it is strongly preferable to rely on direct rather than multi-step iterative forecasting methods. We also find that models estimated with panel data techniques perform only marginally better than those based on individual currency pairs. This finding has bittersweet implications. On the negative side, estimated models encounter a second formidable competitor that, like the RW, bypasses the estimation error problem. On the positive side, the HL model is more acceptable than the RW from the perspective of economic theory.

4. This analysis highlights also that equilibrium exchange rate analysis matters. Simple measures of exchange rate disequilibria, not only signal economic imbalances, but also provide hints in which direction the exchange rate will go.

Our paper has an important message for policymakers. For advanced countries, it is better to rely on the concept of long-run PPP rather than on the RW.

1 Introduction

Not for the first time in the history of the exchange rate literature there is a clear dichotomy between the “in” and “out” of sample evidence. Comprehensive writings have shown that the most popular exchange rate models of our times, albeit successful in explaining what drives them in-sample, cannot consistently beat the random walk (RW) out-of-sample (Cheung et al., 2005, 2017). Building on the results of Ca’ Zorzi et al. (2016, 2017), we find that there are two empirical regularities helping us to beat the RW in a forecasting race. The first one is that Purchasing Power Parity (PPP) holds over the long run. The second one is that in flexible regimes the nominal exchange rate (NER) drives most of the real exchange rate (RER) adjustment process. This evidence is re-assuring as these regularities feature also in the classic Dornbusch (1976) model as well as in state-of-the-art DSGE models (Eichenbaum et al., 2017). These results immediately prompt the question: why did not previous analyses exploit these robust features in the data and beat the RW?

The principal contribution of this paper is to provide an exhaustive and, in our eyes conclusive, answer to this apparent dichotomy. This is achieved in three steps. First, we present robust evidence for the aforementioned two regularities for major bilateral currency pairs of the US dollar. Second, we explain why previous studies, which relied on estimated models, could not systematically beat the RW in light of the pervasive role of the forecast error attributed to estimation. Third, we show that calibrating the half-life RER adjustment to three years and assuming a RW for relative price indices (RPI) is, at least for advanced countries, a simpler and generally better option than forecasting the NER with the RW or relying on estimated models. Direct forecasting or panel data techniques are helpful but fail to beat this simple calibrated model. The beauty of this result is that our approach is so simple that it can be even implemented on the back of a napkin.

2 In-sample regularities on the FX markets

From the IMF-IFS and BIS databases we have taken monthly end-of-period NER against the USD and consumer price index (CPI) data over the period 1975:1-2017:5 for the following countries: Australia (AUD), Canada (CAD), Japan (JPY), New Zealand (NZD), Switzerland (CHF), the United Kingdom (GBP), the euro area (EUR), Korea (KRW), Norway (NOK), Sweden, (SEK) and the United States (US).¹ Using these times series, we have calculated

¹For the euro area over the pre-monetary union period we have taken either a composite of the eleven legacy currencies of the euro (EA11) or Germany (DE). For ease of exposition we report only the results for the EA11 composite measure since the alternative set of results are almost identical.

the bilateral RERs as:

$$rer = ner + rpi, \tag{1}$$

where $rpi = p - p^*$ is the relative price index (domestic vs. US) and ner is the spot nominal exchange rate (USDs per unit of domestic currency) and all variables expressed in logs. Following this definition, an increase in the RER and NER represents an appreciation of the domestic currency with respect to the USD.

The first regularity in the data is that RERs are mean reverting over medium-term horizons. A particularly neat way to illustrate this is to scatter plot changes of real bilateral exchange rate of the euro at different horizons relative to its starting level (top panel, Figure 1). The negative correlation, already visible at the six-month, gets progressively stronger at longer horizons, proving that there is a powerful self-adjusting mechanism at play. The second regularity is illustrated by the middle and bottom panels of Figure 1, which show very clearly how the NER and not the RPI drives this adjustment process. This stylized fact is entirely intuitive, if we think that NERs play an important role in absorbing atypical movements in price competitiveness. This empirical regularity has recently been emphasized by Eichenbaum et al. (2017) and compared to the properties of DSGE models to validate them. However, to be fair, it matches perfectly also one of the standard equations in the Dornbusch model and hence equally validates the open-economy models of the 1970s and 1980s.

Particularly remarkable is how robust these results are to all currency pairs in our dataset. To show this we have estimated the following regressions:

$$\Delta rer_{t,h} = \alpha_{0h} + \alpha_{1h}rer_{t-h} + \epsilon_t \tag{2a}$$

$$\Delta ner_{t,h} = \beta_{0h} + \beta_{1h}\Delta rer_{t,h} + \epsilon_t \tag{2b}$$

$$\Delta rpi_{t,h} = \gamma_{0h} + \gamma_{1h}\Delta rer_{t,h} + \epsilon_t, \tag{2c}$$

where for a variable y we define $\Delta y_{t,h} = y_t - y_{t-h}$. If the RER is mean reverting at long-horizons, then α_{1h} should converge to -1 . Additionally, if the adjustment in RER is driven by NER, rather than RPI, then $\beta_{1h} = 1$ and $\gamma_{1h} = 0$. This is exactly what we find in the data for all currency pairs (Table 1). The same results are confirmed by running a panel regression with “fixed effects”. Full-sample estimates of α_{1h} suggest that the RER adjustment toward PPP is on average completed by 12% after 6 months, 52% after 2 years and more than fully after 5 years. As regards panel data based estimates for β_{1h} and γ_{1h} , they are very close to unity and zero for all horizons h .

3 Out-of-sample evidence

The above in-sample analysis suggests that real and nominal exchange rates do not behave like random walks. But is this assessment compatible with the out-of-sample evidence that most people have in mind? In this section we will assess the accuracy of the forecasts or the RER, NER and RPI generated by a battery of simple models in comparison to that of a RW benchmark. The relative performance of these models is evaluated with the root mean square forecast error (RMSFE) statistics complemented with asymptotic Diebold-Mariano or Clark-West tests. In this paper all the forecasts are generated using rolling regressions with a window of 15 years. Our accuracy measures are hence calculated using errors for forecasts generated from the period 1990:1 onwards. This means that we have 329 one-month-ahead forecasts, 328 two-month-ahead forecasts and so forth. As is standard in the forecasting literature, for each model we report the RMSFE statistics relative to the same statistics for the RW and numbers below one indicate a model that beats the RW.

3.1 Real exchange rate

For the RER we will consider four mean reverting models. The first two models are autoregressive models of order one:

$$rer_t = \mu + \rho(rer_{t-1} - \mu) + \epsilon_t \quad (3)$$

with the only difference that, in one case, the parameters are estimated (AR model), and in the other, calibrated (half-life, HL model). Following the meta-analysis of studies on RER half-life by Rogoff (1996),² we set ρ to a value consistent with a half-life at 3 years, while for μ we take the rolling sample average of the *rer*. As discussed by Ca' Zorzi et al. (2016), if (3) is the true data generating process and ρ is not very distant from unity, it is usually better to forecast with a calibrated HL model than with an AR model, as the impact of estimation error tends to be much more severe than that of misspecification.

The next two competitors are based on regressions presented in model (2a), in which the parameters capturing the pace of adjustment α_{1h} are estimated independently for each horizon h . This could be advantageous to the extent that such methods are less prone to estimation error than the AR iterative approaches and may exploit some non-linearities in the data. The two models differ on how the parameters α_{0h} and α_{1h} are estimated. In the first case these estimates are based on individual time series regression for each bilateral RER (direct forecasts, DF model), whereas in the second case they are based on a panel regression

²Please notice that our calibration is based on studies that were available before the start of the forecast evaluation sample.

with “fixed effects” (panel DF, PDF model). The inclusion of panel data regressions in our horse race is motivated by the desire to estimate α_{1h} more precisely, as suggested by Mark and Sul (2012).

Before turning to forecast evaluation, let us have a look at the set of forecasts derived with the four competing methods using a particular metric, i.e. the pace at which any RER deviation from its recursive mean is absorbed (“PPP absorption” rate). Figure 2 presents the rate of PPP absorption predicted by the four models for the euro-dollar exchange rate from 1990 onwards. A common characteristic across all the models is that at greater horizons the degree of PPP absorption rises. For the direct models this percentage is calculated as $-100 \times \alpha_{h1}$ and for the AR and HL models as $100 \times (1 - \rho^h)$. At the horizon of one-month all models forecast an average absorption rate of about 2%. For all estimated models, including panel data methods, however this number fluctuates sizably pointing to the large role of estimation error. Particularly interesting is also that at longer horizons the direct methods forecast a much higher rate of PPP absorption. For example, at horizons of two years (i.e. $H=24$), the AR and HL models suggest an absorption rate of about 40% while the DF and PDF models of about 60%. The key question is how these differences influence the precision of the RER forecasts.

The outcome of the forecasting competition is presented in Table 2. The main findings in terms of RER forecasting are fourfold. First, in line with Ca’ Zorzi et al. (2016) results the AR model loses against the RW at both short and medium term horizons. Only at horizons of at least five years the mean reverting forces are sufficiently strong to flip the result in favor of the AR model. Second, the DF model, which exploits the regularities reported in the top panels of Figure 1 and Table 1, outforecast the RW at horizons greater than 2 years. This highlights that in this context the “estimation error” problem become less acute with techniques based on direct forecasting relative to iterative methods. Third, extending the analysis to panel data (PDF model), the accuracy of our forecast improves further. However, consistently with the evidence reported by Mark and Sul (2012), these gains are, at least relative to the DF model, marginal. This leads us to our fourth and final finding, i.e. that the calibrated HL model still outperforms all other methods in RER forecasting. We will see later that the final objective of our analysis, i.e. to derive a “good” NER forecast is within easy reach.

3.2 Relative price index

If the RER is forecastable so should be the NER, if we can reasonably extrapolate what drives the RPI. This is tautological if we think in terms of the identity:

$$\Delta ner_{t,h} = \Delta rer_{t,h} - \Delta rpi_{t,h}. \quad (4)$$

But is forecasting relative price indices really easy? The first impression can be deceiving. For example the euro area has shown, for several consecutive years, a tendency to record lower inflation rates than the US. While the direction of the movement has been almost always the same, this (relative) disinflationary process has decelerated in a way that was not easy to anticipate ex-ante.

Let us explore this issue in a more formal setting. Our benchmark is again the RW, which assumes constant RPI over the forecast horizon. This simple approach could be motivated by the importance of global inflation in determining domestic inflation, as suggested by Ciccarelli and Mojon (2010). The first alternative that we propose is to assume that RPI follows an autoregressive process of order one (AR model) with the clear intention to capture different inflation trends across countries and/or some persistence in past inflation rate differentials:

$$\Delta rpi_t = \mu_\pi + \rho_\pi(\Delta rpi_{t-1} - \mu_\pi) + \epsilon_t. \quad (5)$$

The next two models allow the possibility that RPI adjusts to restore equilibrium in the exchange rate market. In this case the forecast is derived from regressions:

$$\Delta rpi_{t,h} = \omega_{0h} + \omega_{1h}rer_{t-h} + \epsilon_t. \quad (6)$$

estimated with time series (DF model) or panel data (PDF model). The last competitor is once again a calibrated half-life model (HL model), in which the parameters from regression (5) are set a priori. In particular, we assume that inflation trends are the same across the two countries by fixing μ_π at 0 and, building on the results of Faust and Wright (2013), ρ is chosen so that half of any inflation differential goes away in six months.

The results presented in the middle panel of Table 3 prove the difficulty to forecast RPI. The AR model extrapolates too much past trends. The DF and PDF models are not that competitive, as the RPI does not play a significant role in the RER adjustment. A marginally better performance than the RW is given by the HL model, as it exploits some short-run persistence of inflation differentials out of sample but the gains are quantitatively negligible.

3.3 Nominal exchange rate

We finally turn to the NER. There are four models that we include in our horse race besides the RW. The first two are based on direct forecasting methods, one estimated with individual time-series (DF model) and the other with panel data (PDF model). In both cases we exploit directly the empirical regularity that the NER adjusts to restore PPP by estimating the following models separately for each horizon h :

$$\Delta ner_{t,h} = \delta_{0h} + \delta_{1h} rer_{t-h} + \epsilon_t \quad (7)$$

The third model, labeled as HL, is in reality an hybrid approach because the RER is forecast with the HL model and RPI with a RW. The fourth model is based on the two half-life models discussed before (3 years for the RER and 6 months for RPI) and labeled for this reason as the 2HL model. All the results are shown in Table 4. The two direct methods (DF and PDF) perform again poorly at short horizons and successfully at horizons greater than 2 years. The two calibrated models (HL and 2HL) are instead extremely competitive at all horizons with the HL model performing marginally better out of the two. The HL model is hence particularly competitive and intuitive. One can also easily calculate the whole forecasting path with this equation:

$$\Delta ner_{t+h,h}^f = \rho^h (rer_t - \overline{rer}), \quad (8)$$

where ρ is calibrated to be consistent with 3-year half-life (close to 0.981 for monthly data), and \overline{rer} is the sample average of the rer .

This finding has bittersweet implications. On the negative side, estimated models encounter a second formidable competitor that, like the RW, bypasses the estimation error problem. On the positive side, the HL model is more acceptable than the RW from the perspective of economic theory and can be implemented easily, even on the back of a napkin, in two steps. Step 1 consists in calculating the initial RER misalignment with an eyeball estimate of what is the distance from the sample mean. Step 2 consists in recalling that, according to this model, one tenth of the required adjustment is achieved by the RER in the first 6 months, one fifth in one year, just over a third in two years and exactly half after 3 years.

Although so easy to compute, such projections are much more accurate than those derived with complex time series models or imposing a constant NER. Simple variants of this approach, by changing the calibration of the half-life adjustment within reasonable values (i.e. between 2 and 5 years), or changing the methodology for calculating the RER equilibrium \overline{rer} would, in general, not change the outcome qualitatively. These variants will similarly

beat the RW by exploiting the mean reversion of the RER while avoiding the common pitfalls described in this paper, i.e. estimation error and poor projections for relatively price indices.

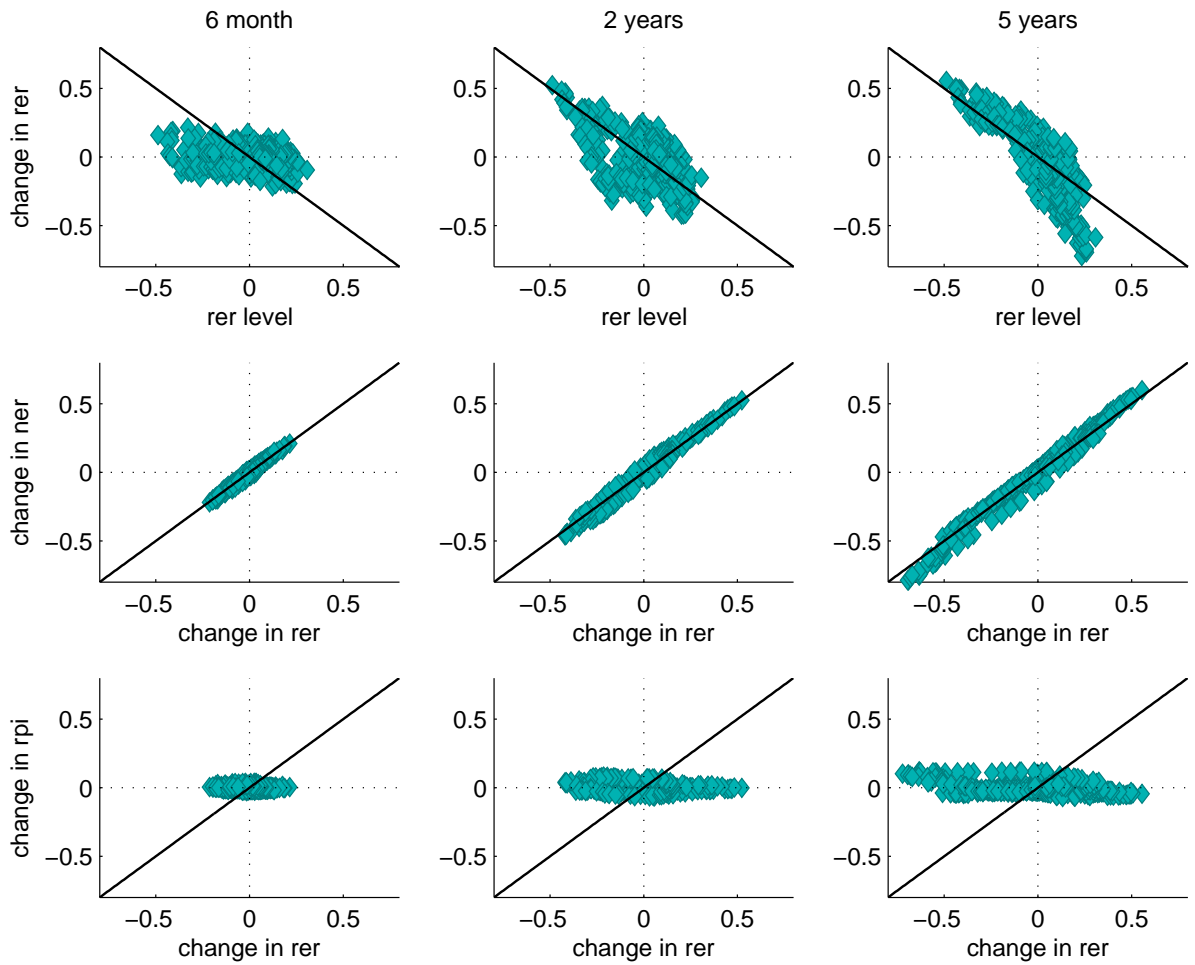
4 Conclusions

Our data suggest that there are two regularities in foreign exchange markets in advanced countries with flexible regimes, i.e. the mean reversion of the RER and the tendency of the NER to drive such adjustment. There are different ways to capitalize on these findings, either with estimated or calibrated models, to forecast exchange rates. The preferable option is to employ a calibrated half-life (HL) model, i.e. to assume a gradual adjustment for the RER and that all the adjustment comes from the NER. Among the estimated approaches, it is clearly better to rely on “direct” rather than “multi-step iterative” forecasting techniques, as they at least outforecast the RW at medium-term horizons. Panel data techniques perform slightly better than those based on individually currency pair but not enough to beat the HL model. The primary intention of this paper was to show how misleading is the common belief that exchange rates are not predictable. It is false that nothing can be said about future movements in exchange rates. They act as shock absorbers. The secret to beat the RW is to impose a reasonable pace at which PPP is restored and assume that relative inflation is zero. This approach is hence simple and yet extremely hard to beat with more sophisticated methods.

References

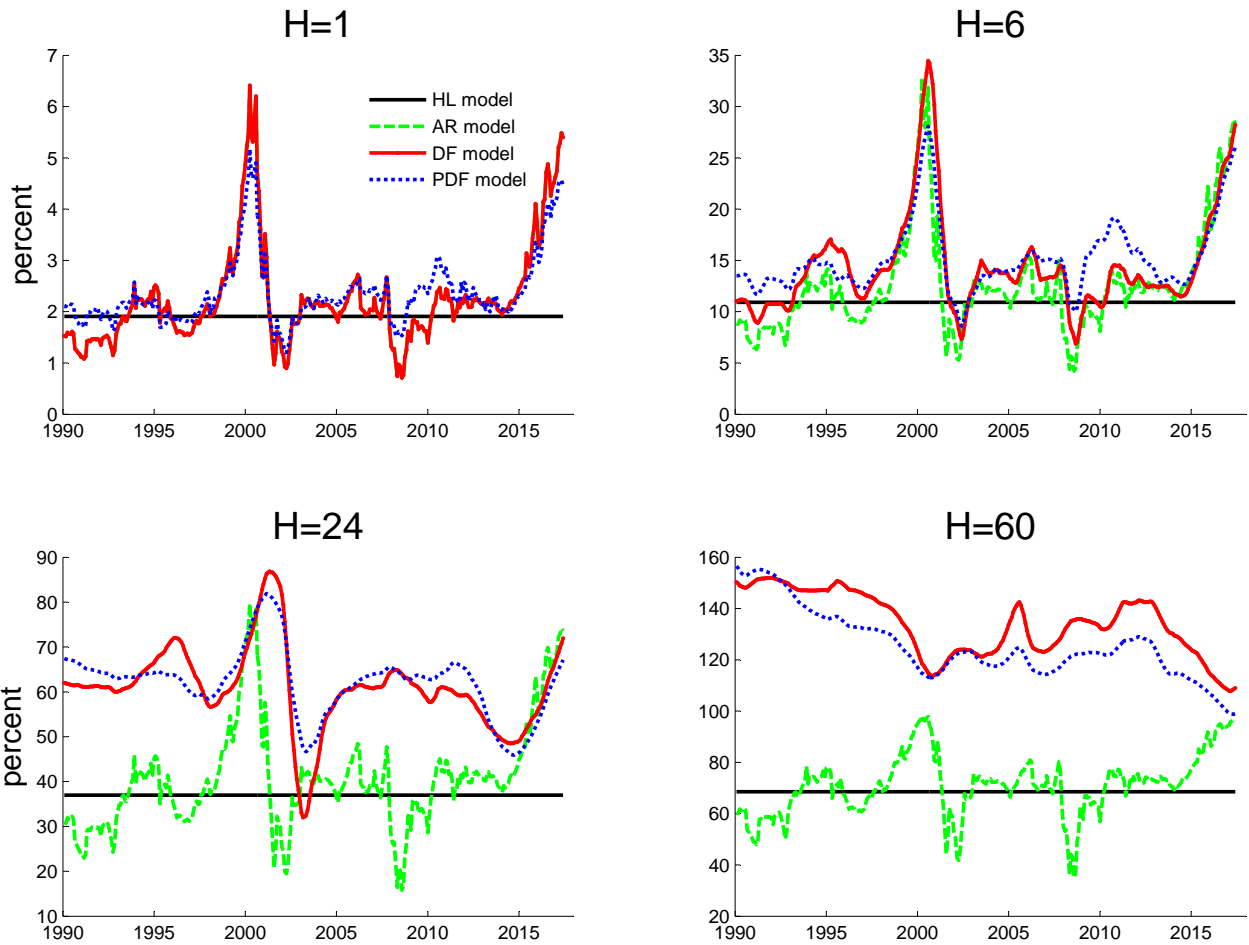
- Ca' Zorzi, M., Kolasa, M., Rubaszek, M., 2017. Exchange rate forecasting with DSGE models. *Journal of International Economics* 107 (C), 127–146.
- Ca' Zorzi, M., Muck, J., Rubaszek, M., 2016. Real exchange rate forecasting and PPP: This time the random walk loses. *Open Economies Review* 27 (3), 585–609.
- Cheung, Y.-W., Chinn, M. D., 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.
- Cheung, Y.-W., Chinn, M. D., Pascual, A. G., Zhang, Y., 2017. Exchange rate prediction redux: new models, new data, new currencies. Working Paper Series 2018, European Central Bank.
- Ciccarelli, M., Mojon, B., 2010. Global inflation. *The Review of Economics and Statistics* 92 (3), 524–535.
- Dornbusch, R., 1976. Expectations and exchange rate dynamics. *Journal of Political Economy* 84 (6), 1161–1176.
- Eichenbaum, M., Johannsen, B. K., Rebelo, S., 2017. Monetary policy and the predictability of nominal exchange rates. NBER Working Papers 23158, National Bureau of Economic Research, Inc.
- Faust, J., Wright, J. H., 2013. Forecasting inflation. In: Elliott, G., Timmermann, A. (Eds.), *Handbook of Economic Forecasting*. Vol. 2. Elsevier, Ch. 1, pp. 2–56.
- Mark, N., Sul, D., 2012. When are pooled panel-data regression forecasts of exchange rates more accurate than the time-series regression forecasts? In: James, J., Marsh, I. W., Sarno, L. (Eds.), *Handbook of Exchange Rates*. John Wiley & Sons, Inc., Ch. 9, pp. 265–281.
- Rogoff, K., 1996. The Purchasing Power Parity puzzle. *Journal of Economic Literature* 34 (2), 647–668.

Figure 1: Exchange rate regularities for the euro dollar



Notes: Changes for a variable Y over horizon H are expressed as $\Delta y_{t,H} = y_t - y_{t-H}$, where $y_t = \log(Y_t)$. The `rer level` is equal to the deviation of rer_{t-H} from the sample mean \overline{rer} .

Figure 2: Rolling absorption rate of RER misalignments of the euro vs. the US dollar



Notes: The plot presents the implied pace of RER adjustment towards the rolling mean. In terms of forecasts, the lines are interpreted as $-100 \times \Delta r_{t+h,h}^f / (r_{er_t} - \bar{r}_{er_t})$, where $\Delta r_{t+h,h}^f$ denotes a h -period ahead forecast elaborated at time t and $r_{er_t} - \bar{r}_{er_t}$ is the percentage RER deviation from the PPP equilibrium estimated at time t .

Table 1: Exchange rate regularities

	6 months			2 years			5 years		
	Estimates of $\Delta rer_{t,h} = \alpha_{0h} + \alpha_{1h}rer_{t-h} + \epsilon_t$								
	α_{0h}	α_{1h}	R^2	α_{0h}	α_{1h}	R^2	α_{0h}	α_{1h}	R^2
AUD	0.00	-0.11	0.06	-0.01	-0.41	0.22	-0.02	-0.92	0.44
CAD	0.00	-0.08	0.04	-0.01	-0.38	0.21	-0.02	-0.98	0.51
JPY	0.00	-0.12	0.06	0.01	-0.52	0.27	0.02	-0.85	0.43
NZD	0.00	-0.11	0.05	0.01	-0.51	0.25	0.00	-1.07	0.48
CHF	0.00	-0.14	0.07	0.01	-0.57	0.29	0.00	-1.05	0.52
GBP	0.00	-0.20	0.09	0.00	-0.84	0.39	0.00	-1.41	0.69
EUR	0.00	-0.13	0.06	-0.01	-0.59	0.29	-0.02	-1.35	0.70
KRW	0.00	-0.17	0.08	0.00	-0.63	0.31	-0.01	-1.16	0.62
NOK	0.00	-0.15	0.07	-0.01	-0.58	0.26	-0.01	-1.36	0.65
SEK	-0.01	-0.09	0.04	-0.02	-0.45	0.21	-0.04	-1.04	0.55
Panel		-0.12	0.06		-0.52	0.26		-1.08	0.54
	Estimates of $\Delta ner_{t,h} = \beta_{0h} + \beta_{1h}\Delta rer_{t,h} + \epsilon_t$								
	β_{0h}	β_{1h}	R^2	β_{0h}	β_{1h}	R^2	β_{0h}	β_{1h}	R^2
AUD	-0.01	1.01	0.96	-0.02	0.99	0.93	-0.04	1.00	0.92
CAD	0.00	1.00	0.96	0.00	0.97	0.94	0.00	0.98	0.95
JPY	0.01	0.97	0.98	0.04	0.98	0.97	0.12	0.99	0.97
NZD	-0.01	0.99	0.94	-0.03	0.91	0.87	-0.08	0.94	0.76
CHF	0.01	1.00	0.98	0.04	0.99	0.95	0.09	0.93	0.91
GBP	0.00	1.02	0.96	-0.01	1.01	0.95	-0.02	1.01	0.93
EUR	0.00	1.02	0.98	0.00	1.03	0.98	0.01	1.05	0.98
KRW	-0.01	1.03	0.94	-0.04	1.04	0.89	-0.09	1.05	0.81
NOK	0.00	1.01	0.96	-0.01	1.01	0.93	-0.02	1.00	0.90
SEK	0.00	1.01	0.97	-0.01	0.99	0.95	-0.02	0.97	0.91
Panel		1.00	0.96		0.99	0.94		0.99	0.90
	Estimates of $\Delta rpi_{t,h} = \gamma_{0h} + \gamma_{1h}\Delta rer_{t,h} + \epsilon_t$								
	γ_{0h}	γ_{1h}	R^2	γ_{0h}	γ_{1h}	R^2	γ_{0h}	γ_{1h}	R^2
AUD	0.01	-0.01	0.00	0.02	0.01	0.00	0.04	0.00	0.00
CAD	0.00	0.00	0.00	0.00	0.03	0.01	0.00	0.02	0.01
JPY	-0.01	0.03	0.05	-0.04	0.02	0.02	-0.12	0.01	0.00
NZD	0.01	0.01	0.00	0.03	0.09	0.06	0.08	0.06	0.01
CHF	-0.01	0.00	0.00	-0.04	0.01	0.00	-0.09	0.07	0.05
GBP	0.00	-0.02	0.01	0.01	-0.01	0.00	0.02	-0.01	0.00
EUR	0.00	-0.02	0.02	0.00	-0.03	0.04	-0.01	-0.05	0.13
KRW	0.01	-0.03	0.01	0.04	-0.04	0.01	0.09	-0.05	0.01
NOK	0.00	-0.01	0.00	0.01	-0.01	0.00	0.02	0.00	0.00
SEK	0.00	-0.01	0.00	0.01	0.01	0.00	0.02	0.03	0.01
Panel		0.00	0.00		0.01	0.00		0.01	0.00

Notes: All regressions are based on monthly data from the period 1975:1-2017:5. The last row represents the estimates of the ‘fixed effect’ panel regressions.

Table 2: RMSFE for the RER with respect to the RW

	AR	DF	PDF	HL	AR	DF	PDF	HL
	1 month				6 months			
AUD	1.01	1.01	1.01	1.00	1.03	1.04	1.05	1.00
CAD	1.01	1.01	1.01	1.00	1.05	1.05	1.08	1.01
JPY	1.01	1.01	1.00	1.00	1.04	1.05	1.02	1.00
NZD	1.01	1.01	1.01	1.00	1.05	1.05	1.05	0.99
CHF	1.01	1.01	1.01	1.00	1.06	1.06	1.04	0.98
GBP	1.01	1.01	1.00	1.00	1.01	1.02	1.00	0.97*
EUR	1.02	1.01	1.01	1.00	1.07	1.07	1.04	0.97
KRW	1.00	1.00	1.00	0.99	1.01	0.98**	0.98*	0.96**
NOK	1.01	1.01	1.01	1.00	1.05	1.05	1.03	0.97
SEK	1.01	1.01	1.01	1.00	1.04	1.05	1.04	0.98
	2 years				5 years			
AUD	1.02	1.11	1.05*	0.96	0.97*	1.08**	1.04**	0.92
CAD	1.07	1.15	1.16*	0.99	1.16	1.15**	0.99**	0.88*
JPY	1.05	1.15	1.02*	0.97	1.01	1.19*	1.23**	0.88*
NZD	1.03	0.96**	0.96**	0.90*	0.87**	0.98**	0.85**	0.81**
CHF	1.04	1.04**	0.98**	0.90*	0.85**	1.01**	0.96**	0.75**
GBP	0.85**	0.76**	0.79**	0.85**	0.66**	0.68**	0.67**	0.71**
EUR	1.03	0.92**	0.92**	0.87**	0.76**	0.67**	0.68**	0.69**
KRW	1.05	0.87**	0.86**	0.86**	1.60	0.79**	0.74**	0.75**
NOK	0.99*	0.93**	0.92**	0.89**	0.78**	0.71**	0.70**	0.73**
SEK	1.05	0.93**	0.91**	0.88*	0.94**	0.91**	0.88**	0.79**

Notes: The table shows the ratio of RMSFE from a given model in comparison to the RMSFE from a RW. Asterisks ** and * denote the 1% and 5% significance levels of the one-sided Diebold-Mariano (HL model) or Clark-West (AR, DF, PDF models) with the alternative that a given model performs better than the RW.

Table 3: RMSFE for the RPI with respect to the RW

	AR	DF	PDF	HL	AR	DF	PDF	HL
	1 month				6 months			
AUD	1.04	1.05	1.03	1.01	1.22	1.39	1.25	1.06
CAD	1.01	1.03	1.03	1.02	1.06	1.19	1.13	1.09
JPY	0.92**	0.94**	0.95**	0.93**	0.76**	0.78**	0.82**	0.76**
NZD	1.19	1.12	1.10	1.02	1.80	1.75	1.70	1.08
CHF	0.97**	1.01**	1.01**	0.96**	0.88**	1.10**	1.04**	0.79**
GBP	1.01	1.01	1.00	1.00	1.05	1.06	1.05	0.99
EUR	1.02	1.00	1.00	1.01	1.01	1.05	1.06	1.08
KRW	0.91**	0.99**	0.96**	0.98	0.91**	0.94**	0.94**	0.95
NOK	1.01	1.02	1.01	1.02	1.10	1.11	1.08	1.04
SEK	1.00	1.02	1.02	0.99	1.09	1.16	1.14	0.94
	2 years				5 years			
AUD	1.71	2.42	1.84	1.07	2.52	3.49	3.00	1.08
CAD	1.21	1.73	1.37	0.93	1.34	2.23	1.60	0.95
JPY	0.58**	0.66**	0.72**	0.77**	0.39**	0.49**	0.46**	0.87**
NZD	3.23	3.71	3.29	1.03	5.00	6.52	5.96	1.03
CHF	0.79**	1.09**	1.00**	0.77**	0.59**	0.79**	0.72**	0.90**
GBP	1.23	1.32	1.29	0.91*	1.57	1.70	1.71	0.96
EUR	1.02**	1.07**	1.02**	1.04	1.12*	1.31	1.24	1.01
KRW	0.78**	0.88**	0.90**	0.87*	0.76**	0.89**	1.00**	0.88**
NOK	1.40	1.66	1.41	1.02	1.92	2.83	2.43	1.05
SEK	1.42	1.49	1.58	0.85**	1.80	2.22	2.17	0.94**

Notes: The table shows the ratio of RMSFE from a given model in comparison to the RMSFE from a RW. Asterisks ** and * denote the 1% and 5% significance levels of the one-sided Diebold-Mariano (HL model) or Clark-West (AR, DF, PDF models) with the alternative that a given model performs better than the RW.

Table 4: RMSFE for the NER with respect to the RW

	DF	PDF	HL	2HL	DF	PDF	HL	2HL
	1 month				6 months			
AUD	1.01	1.01	1.00	1.01	1.04	1.05	0.99	1.01
CAD	1.01	1.01	1.00	1.01	1.05	1.07	1.01	1.03
JPY	1.01	1.00	1.00	1.00	1.05	1.01	1.01	1.00
NZD	1.01	1.01	1.00	1.00	1.04	1.05	0.99	1.01
CHF	1.01	1.00	1.00	1.00	1.04	1.03	0.99	0.99
GBP	1.01	1.01	1.00	1.00	1.03	1.01	0.97*	0.99
EUR	1.02	1.01	1.00	1.00	1.07	1.04	0.97	0.99
KRW	1.00	1.00*	0.99	1.00	0.97**	0.98**	0.97**	0.97*
NOK	1.01	1.01	1.00	1.00	1.05	1.03	0.98	0.98
SEK	1.01	1.01	1.00	1.00	1.05	1.03	0.97	1.00
	2 years				5 years			
AUD	1.05	1.01**	0.95	0.97	0.99**	0.99**	0.89*	0.90*
CAD	1.08	1.11*	1.00	1.01	0.94**	0.88**	0.87*	0.87*
JPY	1.12*	0.98**	0.97	0.98	1.08**	1.12**	0.89	0.88*
NZD	0.98**	0.98**	0.91*	0.92	1.28*	1.14**	0.80**	0.81**
CHF	0.93**	0.90**	0.94	0.93	0.80**	0.78**	0.87*	0.85**
GBP	0.78**	0.81**	0.85**	0.84**	0.86**	0.84**	0.75**	0.74**
EUR	0.91**	0.90**	0.87**	0.88**	0.68**	0.67**	0.68**	0.68**
KRW	0.85**	0.83**	0.88**	0.88**	0.63**	0.58**	0.80**	0.78**
NOK	0.93**	0.92**	0.91**	0.91**	0.77**	0.77**	0.75**	0.76**
SEK	0.88**	0.85**	0.86**	0.88*	0.67**	0.67**	0.69**	0.71**

Notes: The table shows the ratio of RMSFE from a given model in comparison to the RMSFE from a RW. Asterisks ** and * denote the 1% and 5% significance levels of the one-sided Diebold-Mariano (HL and 2HL models) or Clark-West (DF and PDF models) with the alternative that a given model performs better than the RW.

Exchange rate forecasting on a napkin

Michele Ca' Zorzi

European Central Bank, Sonnemannstrasse 20, 60314 Frankfurt, Germany.

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

Michał Rubaszek

SGH Warsaw School of Economics, Al. Niepodległości 162, 02-554 Warsaw, Poland

E-mail: michal.rubaszek@sgh.waw.pl

The views expressed in this paper are those of the authors and not necessarily those of the European Central Bank. We would like to thank, without implicating them, Menzie Chinn, Cedric Tille, Marek Jarocinski and Arnaud Mehl for their excellent comments and suggestions. We also thank Adam Cap for his help in collecting the data.