



WP 05_14

Jiahan Li

University of Notre Dame, USA

Ilias Tsiakas

University of Guelph, Canada

The Rimini Centre for Economic Analysis (RCEA), Italy

Wei Wang

Fifth Third Bank, USA

PREDICTING EXCHANGE RATES OUT OF SAMPLE: CAN ECONOMIC FUNDAMENTALS BEAT THE RANDOM WALK?

Copyright belongs to the author. Small sections of the text, not exceeding three paragraphs, can be used provided proper acknowledgement is given.

The *Rimini Centre for Economic Analysis* (RCEA) was established in March 2007. RCEA is a private, nonprofit organization dedicated to independent research in Applied and Theoretical Economics and related fields. RCEA organizes seminars and workshops, sponsors a general interest journal *The Review of Economic Analysis*, and organizes a biennial conference: *The Rimini Conference in Economics and Finance* (RCEF). The RCEA has a Canadian branch: *The Rimini Centre for Economic Analysis in Canada* (RCEA-Canada). Scientific work contributed by the RCEA Scholars is published in the RCEA Working Papers and Professional Report series.

The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the Rimini Centre for Economic Analysis.

The Rimini Centre for Economic Analysis

Legal address: Via Angherà, 22 – Head office: Via Patara, 3 - 47900 Rimini (RN) – Italy

www.rcfea.org - secretary@rcfea.org

Predicting Exchange Rates Out of Sample: Can Economic Fundamentals Beat the Random Walk?*

Jiahan Li	Ilias Tsiakas	Wei Wang
University of Notre Dame	University of Guelph	Fifth Third Bank
jiahan.li@nd.edu	itsiakas@uoguelph.ca	wei.wang@53.com

February 2014

Abstract

This paper shows that economic fundamentals can generate reliable out-of-sample forecasts for exchange rates when prediction is based on a “kitchen-sink” regression that incorporates multiple predictors. The key to establishing predictability is estimating the kitchen-sink regression with the elastic-net shrinkage method, which improves performance by reducing the effect of less informative predictors in out-of-sample forecasting. Using statistical and economic measures of predictability, we show that our approach outperforms alternative models, including the random walk, individual exchange rate models, a kitchen-sink regression estimated with ordinary least squares, standard forecast combinations and popular ad-hoc strategies such as momentum and the 1/N strategy.

Keywords: Exchange Rates; Out-of-Sample Forecasting; Elastic Net; Combined Forecasts.

JEL Classification: F31; F37; G11; G15; G17.

*This paper is forthcoming in the *Journal of Financial Econometrics*. The authors are grateful to Eric Ghysels (editor), two anonymous referees, Nelson Mark and Lucio Sarno for constructive comments. *Corresponding author:* Ilias Tsiakas, Department of Economics and Finance, University of Guelph, Ontario N1G 2W1, Canada. Tel: 519-824-4120 ext. 53054. Email: itsiakas@uoguelph.ca.

Introduction

Can economic fundamentals generate reliable out-of-sample forecasts for exchange rates? For thirty years, this has been a prominent question in international finance research, but the empirical results are disappointing. Traditional economic models of exchange rate determination have had limited success in explaining and predicting currency movements. As a result, exchange rates are thought to be largely disconnected from economic fundamentals in what is widely known as the “exchange rate disconnect” puzzle (e.g., Engel, Mark and West, 2007). Therefore, the prevailing view in the foreign exchange (FX) literature is that exchange rates are not predictable since models that condition on economic fundamentals cannot outperform a naive random walk model (Meese and Rogoff, 1983).¹

The FX literature typically evaluates the predictive ability of empirical exchange rate models that condition on a particular economic fundamental (e.g., interest rates). However, we can also generate forecasts based on a combination of all available macroeconomic variables. There are two ways to do so. First, we can insert all macroeconomic variables into a single predictive regression. This is known as the “kitchen-sink” regression, which allows multiple predictors to have an effect on the value of a single forecast (see, for example, Welch and Goyal, 2008, for an application on stock returns).² Second, we can estimate several predictive regressions, each capturing the effect of a particular macroeconomic variable, and then combine the several individual forecasts into one forecast combination (see, for example, Rapach, Strauss and Zhou, 2010, also for an application on stock returns). Although the inclusion of multiple sources of predictive information into one forecast seems intuitively attractive, neither kitchen-sink regressions nor forecast combinations have been successful in predicting currency movements (e.g., Della Corte and Tsiakas, 2012).

This paper shows that there is an efficient way to form kitchen-sink forecasts leading to clear statistical and economic evidence of exchange rate predictability. Our empirical analysis uses monthly data on the G10 currencies for a sample period ranging from January 1976 to June 2012, and adopts the following methodology. We first use a number of predictors suggested by popular empirical exchange rate models: uncovered interest parity, purchasing power parity, monetary fundamentals and the Taylor rule. These predictors are inserted into a single predictive regression, the kitchen-sink regression. The key to our approach is that,

¹The exchange rate disconnect puzzle refers primarily to short and medium horizons of up to one year. For longer horizons, there is some evidence that economic fundamentals and exchange rates move together (e.g., Mark, 1995; and Mark and Sul, 2001). For a review of the literature on exchange rate predictability see Della Corte and Tsiakas (2012) and Rossi (2013).

²This is known as the “kitchen-sink” regression simply because it throws “everything but the kitchen sink” into the regression.

instead of ordinary least squares (OLS), we use a shrinkage estimator based on penalized least squares (PLS). The problem with the OLS estimator is that, although it is unbiased, it can have high variance leading to a low degree of predictive accuracy characterized by a high mean squared error (MSE) for the forecasts.

In this paper, we solve this problem by adopting a forecasting methodology known in statistics as the elastic net (Zou and Hastie, 2005), which combines the benefits of ridge regression (Hoerl and Kennard, 1970) and lasso regression (Tibshirani, 1996).³ The elastic net is a shrinkage estimator that imposes two sets of constraints, which shrink all parameter estimates towards zero, while penalizing the smaller and larger parameters more. In other words, the elastic net shrinks the parameter estimates towards the value implied by the benchmark random walk model (i.e., zero). It is important to note that this is achieved in a way that reduces the MSE of both the parameter estimator and of the forecast errors of the model. The application of the elastic-net shrinkage estimator to a kitchen-sink regression for exchange rates leads to a new framework that we term the “efficient kitchen-sink regression.” This framework utilizes efficiently the predictive information from several economic fundamentals by reducing the effect of less informative predictors for exchange rates. As a result, the efficient kitchen-sink regression has substantially higher predictive accuracy than the plain kitchen-sink regression estimated with OLS.

We employ a statistical methodology for evaluating exchange rate predictability that involves testing the null hypothesis of equal predictive ability between the random walk “no-predictability” benchmark and an alternative model, notably the efficient kitchen-sink regression. Following Campbell and Thompson (2008) and Welch and Goyal (2008), our statistical analysis relies heavily on the out-of-sample R^2 measure, R_{oos}^2 . The R_{oos}^2 statistic is based on the out-of-sample MSE of the forecasts generated by the models using either recursive or rolling regressions. We assess the statistical significance of the R_{oos}^2 using the Clark and West (2006, 2007) and the Giacomini and White (2006) t -statistics. We also formally test for the stability of the forecasting ability of the models using the Giacomini and Rossi (2009) t -statistic.

In addition to the statistical evaluation, we also assess the economic value of exchange rate predictability in the context of dynamic asset allocation. A purely statistical analysis is not particularly informative to an investor as it falls short of measuring the tangible economic gains from predictability in exchange rates. Following Della Corte, Sarno and Tsiakas (2009, 2011),

³The idea of the elastic net is to stretch the fishing net so that it retains all the “big fish,” which in our case are the important predictors. The elastic net is a popular method in statistics and is becoming increasingly so in financial econometrics. For an application of the elastic net, see, for example, Bai and Ng (2008), Eickmeier and Ng (2011), Korobilis (2013), Rapach, Strauss and Zhou (2013) and Kim and Swanson (2014).

among others, we design an international asset allocation strategy that exposes a US investor purely to FX risk. We then evaluate the performance of dynamically rebalanced portfolios based on one-month ahead forecasts generated by the empirical models we estimate. We use mean-variance analysis, which allows us to measure how much a risk-averse investor is willing to pay for switching from a portfolio strategy based on the random walk benchmark to the efficient kitchen-sink model. Note that in contrast to univariate statistical measures of forecast accuracy, which are computed separately for each exchange rate, economic value is by design multivariate because it is assessed for the portfolio generated by forecasts on all exchange rate returns. This contributes to the well-known finding that even modest statistical significance in out-of-sample predictive regressions can lead to large economic benefits for investors (see, for example, Campbell and Thompson, 2008).

Using both statistical and economic methods for assessing out-of-sample predictability, we find that the efficient kitchen-sink regression performs better than any alternative. The alternatives include: (i) the random walk benchmark; (ii) four individual empirical exchange rate models: uncovered interest parity, purchasing power parity, monetary fundamentals and the Taylor rule; (iii) the plain kitchen-sink regression estimated with OLS; (iv) the kitchen-sink regression estimated with the simpler shrinkage estimators of ridge regression and lasso regression; and (v) standard forecast combinations typically suggested in the literature, such as combinations based on the (discounted) MSE of each individual forecasting model.

In particular, we find that the efficient kitchen-sink regression is the only model that delivers a positive R_{oos}^2 for all nine exchange rates, which is statistically significant for the majority of the cases. This is true for both recursive and rolling regressions. From an economic point of view, the efficient kitchen-sink regression has a Sharpe ratio net of transaction costs of 0.79 compared to 0.38 for the random walk, where the difference in the Sharpe ratios is statistically significant. Moreover, a risk-averse investor is willing to pay more than 400 basis points every year to be able to use the efficient kitchen-sink forecasts rather than the random walk.

It is also worth noting that the efficient kitchen-sink regression considerably outperforms the two trading strategies of carry trade and momentum that have been very popular in active currency management. The carry trade is a strategy that invests in high-interest currencies by borrowing in low-interest currencies (e.g., Menkhoff, Sarno, Schmeling and Schrimpf, 2012a). This strategy is typically thought to be analogous to the random walk: if exchange rates follow a naive random walk, then the exchange rate return is on average equal to zero, and hence the carry trade expected return is exactly equal to the interest rate differential, thus leading to predictable profits. In contrast, the momentum strategy goes long on currencies that have

recently performed well (the “winners”) and short on currencies that have recently underperformed (the “losers”) (e.g., Menkhoff, Sarno, Schmeling and Schrimpf, 2012b). Finally, the efficient kitchen-sink regression also outperforms the simple 1/N strategy of DeMiguel, Garlappi, and Uppal (2009). This strategy invests in all assets with equal weights and hence involves no forecasting of the asset returns and no optimization in generating the weights in the asset allocation.

The remainder of the paper is organized as follows. In the next section we describe the efficient kitchen-sink methodology and the predictive regressors based on four popular empirical exchange rate models. Section 2 reviews standard forecast combination methods and Section 3 discusses the statistical methodology we use for evaluating exchange rate predictability. In Section 4 we present the dynamic asset allocation framework for assessing the economic value of exchange rate predictability. Section 5 reports our empirical results while Section 6 discusses robustness and extensions. Finally, Section 7 concludes.

1 A Model for Predicting Exchange Rates

1.1 The Kitchen-Sink Regression

Following Welch and Goyal (2008), a kitchen-sink (KS) regression is a term used to describe a single predictive regression that conditions on a large set of predictive variables. The KS regression for forecasting exchange rate returns has the following linear structure:

$$\Delta s_{t+1} = \alpha + \sum_{j=1}^K \beta_j x_{j,t} + \varepsilon_{t+1}, \quad (1)$$

where s_{t+1} is the nominal US dollar spot exchange rate for a particular currency at time $t + 1$, $\Delta s_{t+1} = s_{t+1} - s_t$ is the log-exchange rate return at time $t + 1$, $x_{j,t}$ is the predictor $j \leq K$, α and $\beta = \{\beta_j\}$ are constant parameters to be estimated, and ε_{t+1} is a normal error term. In most empirical applications, the parameters α and β are estimated with OLS. In exchange rate prediction, a typical kitchen-sink regression may incorporate predictors based on uncovered interest parity, purchasing power parity, monetary fundamentals and the Taylor rule. In a following section, we discuss these predictors in more detail.

Although intuitively it seems attractive to utilize all available information in a single KS regression, in practice the out-of-sample performance of the KS regression estimated with OLS has been abysmal (see, for example, Welch and Goyal, 2008, for stock returns). Our empirical analysis explores the role of shrinkage estimation in improving the predictive accuracy of the KS regression. We focus on a particular estimator, the elastic net, which as we will show

empirically will transform the KS regression from being the worst performing model when estimated with OLS to the best performing model when estimated with the elastic net.

1.1.1 The Role of Shrinkage Estimation

We begin by adopting the mean squared error (MSE) as the loss function with which we evaluate the statistical performance of a set of forecasting models for exchange rates. Hence our primary statistical objective is to find the model with the lowest out-of-sample MSE. In choosing an estimation method, OLS appears attractive because among unbiased estimators it has the lowest variance and hence the lowest MSE. However, by relaxing the condition of unbiasedness, we may be able to considerably improve the MSE of a model.

By design, shrinkage estimation provides a biased estimator which, however, may have lower variance and lower MSE than OLS. In other words, a shrinkage estimator may be more efficient than OLS because it improves the variance-bias tradeoff and may lead to more accurate forecasts, which is critically important in empirical applications. As a result, shrinkage methods have been popular in statistics (e.g., Tibshirani, 1996), econometrics (e.g., Bai and Ng, 2008; De Mol, Giannone and Reichlin, 2008; and Stock and Watson, 2012), and portfolio management (e.g., Jorion, 1985, 1986, 1991; Ledoit and Wolf, 2003, 2004; Jagannathan and Ma, 2003; DeMiguel, Garlappi, Nogales and Uppal, 2009; and Rapach, Strauss and Zhou, 2013). For these reasons, our empirical analysis implements shrinkage estimation of the KS regression. This will shrink the regression coefficients towards zero, which is the value implied by the benchmark random walk model.⁴

To further motivate the role of shrinkage consider the following example. Suppose that according to the true model, which is unknown to us, economic fundamentals are somehow related to future exchange rate movements. Then, by setting the regression coefficients to zero, the random walk model produces biased estimates with zero variance. In contrast, OLS estimation of the KS regression that conditions on all economic fundamentals will produce unbiased estimates with potentially high variance. Then, a shrinkage estimator of the KS regression can be thought of as a combination of the two cases, which can lead to a biased estimator with low variance. In this case, Tu and Zhou (2011) provide an important economic

⁴Consider the simple linear regression: $y = X\beta + \varepsilon$, where $E[\varepsilon] = 0$ and $Var[\varepsilon] = \sigma^2$. The MSE of the β estimator is defined as $MSE(\hat{\beta}) = E[(\hat{\beta} - \beta)]^2 = [Bias(\hat{\beta})]^2 + Var(\hat{\beta})$, where $Bias(\hat{\beta}) = E[\hat{\beta}] - \beta$. The OLS estimator is unbiased and hence $Bias(\hat{\beta}) = 0$. The shrinkage estimator leads to a biased estimator, which however can have substantially lower variance than OLS so that the MSE of the shrinkage estimator is lower than that of the OLS estimator. Note that the MSE for the forecasts of the model is directly related to the MSE of the estimator since it can be shown that: $MSE(y - \hat{y}) = E[(y - \hat{y})]^2 = MSE(\hat{\beta}) + \sigma^2$ (see Tibshirani, 1996). By reducing $MSE(\hat{\beta})$, the shrinkage estimator may lead to a considerable reduction in the MSE of the forecasts and a marked improvement in the model's predictive accuracy.

insight: "...a concave utility investor will prefer a suitable average of good and bad performances to either a good or a bad performance randomly..."⁵ Hence a combination in the form of a shrinkage estimator of the kitchen-sink regression that shrinks parameters towards zero will be preferred to either the random walk, where the parameters are fixed at zero, or the plain kitchen-sink regression where the parameters are estimated with OLS. As a result, although shrinkage may push all coefficients close to zero, it may still generate forecasts which are better than the benchmark random walk forecasts.

1.1.2 The Role of the Elastic Net

The main aspect of our approach is that we shrink the regression coefficients in a particular way by estimating the KS regression with the elastic-net methodology of Zou and Hastie (2005). The elastic net is a shrinkage estimator based on penalized least squares (PLS) that is popular in statistics and solves the following system:

$$\begin{aligned} & \min_{\beta} \frac{1}{2} \sum_{t=1}^T \left(\Delta s_{t+1} - \alpha - \sum_{j=1}^K \beta_j x_{j,t} \right)^2 \\ \text{s.t.} & \quad \sum_{j=1}^K |\beta_j| < s_1 \\ \text{and} & \quad \sum_{j=1}^K \beta_j^2 < s_2, \end{aligned} \tag{2}$$

where s_1 and s_2 are positive constants to be estimated. We refer to this model as the efficient kitchen-sink (E-KS) regression because it potentially leads to a shrinkage kitchen-sink estimator with a substantially improved variance-bias tradeoff and a lower MSE, something that will be shown empirically. The elastic net works as follows.⁶

When $s_1 = \infty$, i.e., the first constraint is unbounded, then Eq. (2) reduces to the ridge regression (Hoerl and Kennard, 1970), denoted as RR. When $s_2 = \infty$, i.e., the second constraint is unbounded, then Eq. (2) is called lasso regression, denoted as LR, which stands for "least absolute shrinkage and selection operator" (Tibshirani, 1996). Both RR and LR are shrinkage estimators, which shrink regression coefficients towards zero. The main difference in these two estimators lies in their shrinkage intensity for small and large coefficients. Ridge regression (with the only constraint $\sum_{j=1}^K \beta_j^2 < s_2$) tends to shrink more the large regression coefficients.

In contrast, lasso regression (with the only constraint $\sum_{j=1}^K |\beta_j| < s_1$) tends to shrink more the

⁵See Tu and Zhou (2011), page 205.

⁶We have also estimated an E-KS model with all the interaction terms between the predictors. However, the performance of this model is virtually identical to the E-KS. Hence we do not report these results but they are available upon request.

small regression coefficients.⁷

Furthermore, ridge regression does not produce a parsimonious model as it keeps all predictors in the model. In contrast, lasso regression sets some coefficients to exactly zero. For example, if a group of predictors are highly correlated, then the lasso regression tends to select only one predictor from the group and does not care which one is selected. In the end, the combination of ridge and lasso regression in the elastic-net regression provides a flexible shrinkage scheme that may lead to more accurate forecasts. The estimation algorithm for E-KS regression using the elastic net is summarized in Appendix A.

In addition to estimating the E-KS model, we also estimate the RR and the LR model. This will allow us to evaluate whether any potential gains come from the RR constraint, the LR constraint or from the combination of the two constraints. More generally, as all three models implement shrinkage estimation, a comparison of their performance will allow us to determine whether it is shrinkage that drives the results or the particular type of shrinkage that combines the two constraints into the E-KS model.

1.2 Regressors Based on Empirical Exchange Rate Models

Our empirical analysis uses a set of four regressors based on standard theories of exchange rate determination. Each regressor j defines the predictive variable $x_{j,t}$, $j = 1, \dots, K = 4$, that is used to forecast exchange rate returns in the context of the efficient kitchen-sink regression.

1.2.1 Uncovered Interest Parity

The first regressor is based on the uncovered interest parity (UIP) condition. UIP holds under risk neutrality and rational expectations, and implies three equivalent statements: (i) the forward rate is an unbiased estimator of the future spot rate; or (ii) the expected exchange rate return is equal to the interest rate differential; or (iii) the expected FX excess return is

⁷To be more precise, note that if all predictors are uncorrelated and standardized, it can be shown that ridge regression shrinks all estimates proportionally (i.e., it scales all estimates by the same constant between 0 and 1), which means that in absolute amount larger estimates shrink more. In contrast, lasso shrinks all estimates by the same amount (i.e., subtracts the same constant from all estimates), which means that proportionally the smaller estimates shrink more.

equal to zero.⁸ The UIP regressor is specified as follows:

$$x_{1,t} = f_t - s_t, \quad (3)$$

where f_t is the log of the one-month forward exchange rate at time t , which is the rate agreed at time t for an exchange of currencies at $t + 1$.

To motivate the UIP regressor, suppose that we estimate a linear predictive regression conditioning only on this regressor. Then, the UIP condition implies that $\alpha = 0$, $\beta = 1$, and the error term is serially uncorrelated. However, empirical studies consistently reject this condition and it is a stylized fact that estimates of β often display a negative sign (e.g., Bilson, 1981; Fama, 1984). The empirical rejection of UIP is widely known as the “forward bias” because it implies that the forward rate is a biased estimator of the future spot rate or, alternatively, that high-interest rate currencies tend to appreciate rather than depreciate over time.

1.2.2 Purchasing Power Parity

The second regressor is based on the purchasing power parity (PPP) condition. PPP is a long-run condition, which states that national price levels should be equal when expressed in a common currency (e.g., Taylor and Taylor, 2004). The PPP regressor is as follows:

$$x_{2,t} = p_t - p_t^* - s_t, \quad (4)$$

where p_t is the log of the domestic price level and p_t^* the log of the foreign price level.

1.2.3 Monetary Fundamentals

The third regressor conditions on monetary fundamentals (MF) as follows:

$$x_{3,t} = (m_t - m_t^*) - (y_t - y_t^*) - s_t, \quad (5)$$

where m_t is the log of the domestic money supply and y_t is the log of the domestic real output. Similarly, m_t^* is the log of the foreign money supply and y_t^* is the log of the foreign real output.

⁸The monthly FX return is defined as $s_{t+1} - s_t$, whereas the monthly FX excess return is defined as $s_{t+1} - f_t = s_{t+1} - s_t - (i_t - i_t^*)$, where f_t is the log of the one-month forward exchange rate at time t , i_t and i_t^* are the domestic and foreign one-month nominal interest rates, respectively. Note that the equality is due to the covered interest parity (CIP). If CIP holds, then the interest rate differential is equal to the forward premium, $f_t - s_t = i_t - i_t^*$. There is ample empirical evidence that CIP holds in practice for the data frequency examined in this paper. For recent evidence, see Akram, Rime and Sarno (2008). The only exception in our sample is the period following Lehman Brothers’ bankruptcy, when a violation of the CIP persisted for a few months (e.g., Mancini-Griffoli and Ranaldo, 2011).

The empirical evidence on the relation between exchange rates and fundamentals is mixed. On the one hand, a large literature has established that short-run exchange rate movements appear to be disconnected from the underlying economic fundamentals (e.g., Engel, Mark and West, 2007). On the other hand, there is evidence that exchange rates and fundamentals move together in the long run (e.g., Mark, 1995; Groen, 2000; Rapach and Wohar, 2002).

1.2.4 Taylor Rule

The final regressor is based on the (asymmetric) Taylor (1993) rule as follows:

$$x_{4,t} = 1.5 (\pi_t - \pi_t^*) + 0.1 (y_t^g - y_t^{*g}) + 0.1 (s_t + p_t^* - p_t), \quad (6)$$

where π_t is the domestic inflation rate, π_t^* the foreign inflation rate, y_t^g the domestic output gap, and y_t^{*g} the foreign output gap. The output gap is measured as the percent deviation of real output from an estimate of its potential level computed using the Hodrick and Prescott (1997) filter.⁹ The parameters on the inflation difference (1.5), output gap difference (0.1) and the real exchange rate (0.1) are fairly standard in the literature (e.g., Engel, Mark and West, 2007; Mark, 2009; Molodtsova and Papell, 2009).¹⁰

1.3 The Random Walk Benchmark

Since the seminal contribution of Meese and Rogoff (1983), the random walk (RW) model has become the benchmark in assessing exchange rate predictability. For this reason, we use the RW model as the universal benchmark against which we compare all models. The RW captures the prevailing view in the FX literature that exchange rates are not predictable when conditioning on economic fundamentals, especially at short horizons. The RW model also forms the basis of the widely used carry trade strategy in active currency management (e.g., Burnside, Eichenbaum, Kleshchelski and Rebelo, 2011; Lustig, Roussanov and Verdelhan, 2011; Menkhoff, Sarno, Schmeling and Schrimpf, 2012a).

Note that if the RW holds, then spot exchange rate returns cannot be predicted by economic fundamentals. In contrast, however, FX excess returns are perfectly predictable and equal to the interest rate differential. The portfolio returns of a dynamic FX strategy depend

⁹We construct the output gap using the Hodrick and Prescott (1997) filter by setting the smoothing parameter to be equal to 14,400 as in Molodtsova and Papell (2009). In estimating the Hodrick-Prescott trend out of sample, at any given period t , we only use data up to period $t - 1$. We then update the trend every time a new observation is added to the sample. This captures as closely as possible the information available at the time a forecast is made and avoids look-ahead bias.

¹⁰We also estimate (rather than fix) the parameters on the inflation difference, output gap difference and the real exchange rate but we find that the results remain qualitatively identical. Hence we use the fixed parameters as above.

on excess returns not just on spot returns. Hence, even if the RW holds, interest rates are important in determining the returns to a dynamic asset allocation strategy for currencies. We provide more details later on what determines portfolio returns.

1.4 Asset Pricing Foundations

The Engel and West (2005) asset pricing model nests and motivates the regressors described above. This model builds on earlier work on pricing stock returns by Campbell and Shiller (1987, 1988) and West (1988). Its main implication is that economic fundamentals and the exchange rate are linked as follows:

$$s_t = (1 - b)(f_{1,t} + z_{1,t}) + b(f_{2,t} + z_{2,t}) + bE_t s_{t+1}, \quad (7)$$

where $f_{i,t}$ ($i = 1, 2$) are the observed economic fundamentals and $z_{i,t}$ are the unobserved fundamentals that drive the exchange rate. Iterating forward and imposing the no-bubbles condition leads to the following present-value relation:

$$s_t = (1 - b) \sum_{j=0}^{\infty} b^j E_t (f_{1,t+j} + z_{1,t+j}) + b \sum_{j=0}^{\infty} b^j E_t (f_{2,t+j} + z_{2,t+j}). \quad (8)$$

Two aspects of this model are worth noting. First, the model motivates the use of the four regressors discussed above either individually or together in a kitchen-sink regression. All these cases take the general form of Equation (7), and we discuss two specific examples in Appendix B. Second, Engel and West (2005) show that, although exchange rates are related to economic fundamentals, they may still appear to follow a random walk if the discount factor b is close to one, and either (a) or (b) hold, where: (a) $f_{1,t} + z_{1,t} \sim I(1)$ and $f_{2,t} + z_{2,t} = 0$; (b) $f_{2,t} + z_{2,t} \sim I(1)$.

2 Combined Forecasts

The E-KS is a single predictive regression that incorporates multiple predictors leading to a single forecast. We now turn to discussing methods that first estimate several predictive regressions each with a single predictor. These are used to generate a set of individual forecasts, which are then combined into a single forecast combination (e.g., Bates and Granger, 1969).¹¹

In particular, for each exchange rate, we estimate $K^* = K + 1$ predictive regressions, which include $K = 4$ regressions that condition on the regressors (UIP, PPP, MF, TR) plus

¹¹For recent applications of combined forecasts see Stock and Watson (2004), Timmermann (2006), Wright (2008), and Rapach, Strauss and Zhou (2010).

the RW.¹² Each predictive regression $j \leq K^*$ generates an individual forecast $\Delta s_{j,t+1|t}$ for the one-month ahead exchange rate returns. We define the combined forecast $\Delta s_{c,t+1|t}$ as the weighted average of the K^* individual forecasts:

$$\Delta s_{c,t+1|t} = \sum_{j=1}^{K^*} \theta_{j,t} \Delta s_{j,t+1|t}, \quad (9)$$

where $\{\theta_{j,t}\}_{j=1}^{K^*}$ are the ex-ante combining weights determined at time t . The objective of this exercise is to assess whether standard methods of combining economic fundamentals lead to reliable exchange rate forecasts and to compare these forecasts to the E-KS model. Note that, like the E-KS model, the forecast combination methods produce ex ante forecasts that can be used in realistic trading strategies.

The combining methods we consider differ in how the weights are determined and can be organized into two sets. The first set uses simple averaging schemes: mean, median, and trimmed mean. The mean combination forecast sets $\theta_{j,t} = 1/K^*$ in Equation (9); the median combination forecast is the median of $\{\Delta s_{j,t+1|t}\}_{j=1}^{K^*}$; and the trimmed mean combination forecast sets $\theta_{j,t} = 0$ for the two individual forecasts with the smallest and largest values, and $\theta_{j,t} = 1/(K^* - 2)$ for the remaining individual forecasts in Equation (9). These combined forecasts disregard the historical performance of the individual forecasts.

The second set of combined forecasts uses statistical information on the past out-of-sample performance of each individual model. In particular, we compute the discounted MSE (*DMSE*) forecast combination by setting the following weights:

$$\theta_{j,t} = \frac{DMSE_{j,t}^{-1}}{\sum_{i=1}^{K^*} DMSE_{i,t}^{-1}}, \quad DMSE_{j,t} = \sum_{t=M+1}^{T-1} \varphi^{T-1-t} (\Delta s_{j,t+1} - \Delta \hat{s}_{j,t+1})^2, \quad (10)$$

where φ is a discount factor and M are the in-sample observations on which we condition to form the first out-of-sample forecast. For $\varphi < 1$, greater weight is attached to the most recent forecasts. The *DMSE* forecasts are computed for three values of $\varphi = \{0.90, 0.95, 1.0\}$. The case of no discounting ($\varphi = 1$) corresponds to the Bates and Granger (1969) forecast combination. We also compute simpler “most recently best” *MSE* (κ) forecast combinations that use no discounting ($\varphi = 1$) and weigh individual forecasts by (the inverse of) the *DMSE* computed over the last κ months, where $\kappa = \{12, 36, 60\}$.

¹²Our forecast combinations include the RW. If we exclude it, the results remain qualitatively the same and are available upon request.

3 Statistical Evaluation of Exchange Rate Predictability

We evaluate the performance of the empirical exchange rate models with statistical tests of equal out-of-sample predictive ability between one of the models we estimate (e.g., E-KS) and the benchmark RW model. In effect, we are comparing the performance of a parsimonious restricted null model (the RW, where $\beta = 0$) to a set of larger alternative unrestricted models that nest the parsimonious model (where $\beta \neq 0$).

We first estimate all the empirical models, and then run an out-of-sample forecasting exercise as follows. Given the full sample of observables $\{\Delta s_{t+1}, x_t\}_{t=1}^{T-1}$, we define an in-sample (IS) period using observations $\{\Delta s_{t+1}, x_t\}_{t=1}^M$, and an out-of-sample (OOS) period using $\{\Delta s_{t+1}, x_t\}_{t=M+1}^{T-1}$. This exercise produces $F = (T - 1) - M$ OOS forecasts. The IS period for x_t ranges from January 1976 to December 1986.¹³ The first OOS forecast is for the February 1987 value of Δs_{t+1} that conditions on the January 1987 value of x_t . The last forecast is for June 2012. Therefore, our analysis uses $T - 1 = 437$ monthly observations, $M = 132$ and $F = 305$. The OOS monthly forecasts are obtained in two ways: (i) with recursive regressions for the period of January 1987 to June 2012 that successively re-estimate the model parameters every time a new observation is added to the sample; and (ii) with rolling regressions using a 10-year window that generate forecasts for the same forecasting period.

In what follows, we describe the main statistical criterion for evaluating the OOS predictive ability of the models: the Campbell and Thompson (2008) and Welch and Goyal (2008) OOS R^2 statistic, R_{oos}^2 . In our application, the R_{oos}^2 compares the unconditional one-month ahead forecasts $\Delta \bar{s}_{t+1|t}$ of the benchmark RW model to the conditional forecasts $\Delta \hat{s}_{t+1|t}$ of an alternative model. Then, the R_{oos}^2 statistic is given by:

$$R_{oos}^2 = 1 - \frac{MSE(\Delta \hat{s}_{t+1|t})}{MSE(\Delta \bar{s}_{t+1|t})} = 1 - \frac{\sum_{t=M+1}^{T-1} (\Delta s_{t+1} - \Delta \hat{s}_{t+1|t})^2}{\sum_{t=M+1}^{T-1} (\Delta s_{t+1} - \Delta \bar{s}_{t+1|t})^2}. \quad (11)$$

A positive R_{oos}^2 implies that the alternative model outperforms the benchmark RW by having a lower MSE.

We assess the statistical significance of the R_{oos}^2 statistic by applying two testing procedures: Clark and West (2006, 2007) and Giacomini and White (2006). Both are tests of the

¹³The IS period is 11 years, which is 10 years plus the one year we need to compute a rolling value of annual inflation.

null hypothesis of equal OOS predictive ability between one of the models we estimate (e.g., E-KS) and the benchmark RW model. We first describe the Clark and West (2006, 2007) procedure, which is particularly useful because it accounts for the fact that, under the null, the MSE of the RW benchmark is expected to be lower. This is because for any of the alternative models we have to estimate a parameter vector that, under the null, is not helpful in prediction, thus introducing noise into the forecasting process. Therefore, finding that the RW has smaller MSE (i.e., negative R_{oos}^2) is not clear evidence against the alternative model. Clark and West (2006, 2007) propose to adjust the MSE as follows:

$$MSE_{adj} = \frac{1}{F} \sum_{t=M+1}^{T-1} (\Delta s_{t+1} - \Delta \hat{s}_{t+1|t})^2 - \frac{1}{F} \sum_{t=M+1}^{T-1} (\Delta \bar{s}_{t+1|t} - \Delta \hat{s}_{t+1|t})^2. \quad (12)$$

Then, a computationally convenient way of testing for equal MSE (i.e., whether R_{oos}^2 is zero) is to define:

$$\widehat{test}_{t+1} = (\Delta s_{t+1} - \Delta \bar{s}_{t+1|t})^2 - [(\Delta s_{t+1} - \Delta \hat{s}_{t+1|t})^2 - (\Delta \bar{s}_{t+1|t} - \Delta \hat{s}_{t+1|t})^2], \quad (13)$$

and to regress \widehat{test}_{t+1} on a constant, using the t -statistic for a zero coefficient. Even though the asymptotic distribution of this test is non-standard (e.g., McCracken, 2007), Clark and West (2006, 2007) show that standard normal critical values provide a good approximation, and therefore recommend to reject the null of equal predictive ability if the statistic is greater than +1.282 (for a one-sided 0.10 test) or +1.645 (for a one-sided 0.05 test) or +2.326 (for a one-sided 0.01 test).¹⁴

The Clark and West (2006, 2007) testing procedure is asymptotic and hence relies on the population values of coefficients. We also assess the OOS predictive ability of the models using the Giacomini and White (2006) test that relies on the more realistic case of using finite-sample estimated coefficients. This is a more general test that effectively evaluates two different forecasting methods, and hence not only does it take into account that the forecasting models are different but also accounts for the estimation procedure used for each model and the choice of estimation window. Furthermore, the test allows for a unified treatment of nested and non-nested models. The null hypothesis of the Giacomini and White (2006) test is that the *estimated* OOS MSE of a model is equal to the *estimated* OOS MSE of the benchmark. The null hypothesis of equal predictive ability is assessed with a two-sided t -statistic.¹⁵

¹⁴Note that the Clark and West (2006, 2007) statistic is testing the null hypothesis of equal predictive accuracy in population, while the reported R_{oos}^2 values reflect finite-sample performance. This implies that a rejection of the null hypothesis may occasionally be associated with a negative R_{oos}^2 .

¹⁵Recall that our empirical analysis uses both rolling and recursive forecasts. Strictly speaking, the asymptotic results of the Giacomini and White (2006) test require that the forecasting models are estimated with a

In addition to the two methods for testing the significance of the R_{oos}^2 , we also use the formal testing procedure of Giacomini and Rossi (2009) to test for the stability of the forecasting ability of the models. This test is designed to detect forecast breakdowns by assessing whether a model that provides good forecasts over one period can continue doing so over a subsequent period. In our context, the null hypothesis of this test is that the OOS MSE of a model is equal to the IS MSE of the same model. We test the null hypothesis with a one-sided t -statistic used for both recursive and rolling forecasts. The one-sided t -test focuses on the alternative that the OOS MSE of model is higher than the IS MSE. Note that we do not formally test the two-sided alternative that also considers whether the OOS MSE of a model is lower than its IS MSE, since this is actually desirable and does not constitute a forecast breakdown.

4 Economic Evaluation of Exchange Rate Predictability

This section describes the framework for providing an economic evaluation of exchange rate predictability based on the dynamic asset allocation methodology of Della Corte, Sarno and Tsiakas (2009, 2011).

4.1 The Dynamic FX Strategy

We design an international asset allocation strategy that involves trading the US dollar and nine other currencies: the Australian dollar, Canadian dollar, Swiss franc, Deutsche mark/euro, British pound, Japanese yen, Norwegian krone, New Zealand dollar and Swedish krona. Consider a US investor who builds a portfolio by allocating her wealth between ten bonds: one domestic (US), and nine foreign bonds (Australia, Canada, Switzerland, Germany, UK, Japan, Norway, New Zealand and Sweden). The yield of the bonds is proxied by Eurodeposit rates. At each period $t + 1$, the foreign bonds yield a riskless return in local currency but a risky return r_{t+1} in US dollars. The expected US dollar return of investing in a foreign bond is equal to $r_{t+1|t} = i_t^* + \Delta s_{t+1|t}$, where $r_{t+1|t} = E_t[r_{t+1}]$ is the conditional expectation of r_{t+1} , i_t^* is the foreign nominal interest rate, and $\Delta s_{t+1|t} = E_t[\Delta s_{t+1}]$ is the conditional expectation of Δs_{t+1} . As the interest rate i_t^* is known in advance at time t , the only risk the US investor is exposed to from time t to $t + 1$ by investing in a foreign bond is FX risk. Similarly, investing in the US bond is riskless, and hence the riskless rate r_f is equal to the yield of the US bond.

Every month the investor takes two steps. First, she uses each model to forecast the rolling estimation scheme and, therefore, do not apply under a recursive estimation scheme. However, Monte Carlo evidence reported in Clark and McCracken (2013) suggests that testing the finite-sample null hypothesis of equal predictive accuracy works well also for recursive forecasts. Note that Clark and McCracken (2013) provide a survey of statistical tests for evaluating forecasts at the population level and in the finite sample.

one-month ahead exchange rate returns. Second, conditional on the forecasts of each model, she dynamically rebalances her portfolio by computing the new optimal weights. This setup is designed to assess the predictive ability of exchange rate forecasts generated by the EKS regression by informing us whether these forecasts lead to a better performing allocation strategy than conditioning on the benchmark RW model.

4.2 Mean-Variance Dynamic Asset Allocation

Mean-variance analysis is a natural framework for evaluating strategies that exploit predictability in the mean and variance. Consider an investor who has a one-month horizon and constructs a dynamically rebalanced portfolio. Computing the time-varying weights of this portfolio requires one-month ahead forecasts of the conditional mean and the conditional variance-covariance matrix. Let r_{t+1} denote the $N \times 1$ vector of risky asset returns at time $t + 1$, and $\Sigma_{t+1|t} = E_t[(r_{t+1} - r_{t+1|t})(r_{t+1} - r_{t+1|t})']$ the $N \times N$ conditional variance-covariance matrix of r_{t+1} .

Our analysis focuses on the maximum expected return strategy, which leads to an allocation on the efficient frontier and is often used in active currency management. This strategy maximizes the expected portfolio return at each month t for a given target portfolio volatility:

$$\begin{aligned} \max_{w_t} \quad & r_{p,t+1|t} = w_t' r_{t+1|t} + (1 - w_t' \iota) r_f \\ \text{s.t.} \quad & \sigma_p^* = (w_t' \Sigma_{t+1|t} w_t)^{1/2}, \end{aligned} \quad (14)$$

where σ_p^* is the target conditional volatility of the portfolio returns, and ι is an $N \times 1$ vector of ones. The solution to the maximum expected return rule gives the following risky asset weights:

$$w_t = \frac{\sigma_p^*}{\sqrt{C_t}} \Sigma_{t+1|t}^{-1} (\mu_{t+1|t} - \iota r_f), \quad (15)$$

where $C_t = (\mu_{t+1|t} - \iota r_f)' \Sigma_{t+1|t}^{-1} (\mu_{t+1|t} - \iota r_f)$.

Then, the return on the investor's portfolio is:

$$r_{p,t+1} = w_t' r_{t+1} + (1 - w_t' \iota) r_f. \quad (16)$$

Note that we assume that $\Sigma_{t+1|t} = \bar{\Sigma}$, where $\bar{\Sigma}$ is the unconditional covariance matrix of exchange rate returns. In other words, we do not model the dynamics of FX return volatility and correlation. Therefore, the optimal weights will vary across the empirical exchange rate models only to the extent that the predictive regressions produce better forecasts of the exchange rate returns.¹⁶

¹⁶See Della Corte, Sarno and Tsiakas (2012) for an economic evaluation of volatility and correlation timing

4.3 Performance Measures

We evaluate the performance of the exchange rate models using the Goetzmann, Ingersoll, Spiegel and Welch (2007) performance measure defined as:

$$M(r_p) = \frac{1}{(1-\gamma)} \ln \left\{ \frac{1}{T} \sum_{t=1}^T \left(\frac{1+r_{p,t}}{1+r_f} \right)^{1-\gamma} \right\}, \quad (17)$$

where γ denotes the investor's degree of relative risk aversion (RRA). $M(r_p)$ can be interpreted as the certainty equivalent of the excess portfolio returns. This is an attractive criterion since it is robust to the distribution of portfolio returns and does not require the assumption of a particular utility function to rank portfolios.

We compare the performance of the E-KS model to the benchmark RW by computing the difference:

$$\mathcal{P} = M(r_p^*) - M(r_p^b), \quad (18)$$

where r_p^* are the portfolio returns of the E-KS strategy and r_p^b are the portfolio returns of the benchmark RW. We interpret \mathcal{P} as the maximum performance fee an investor will pay to switch from the RW to the E-KS strategy. In other words, this performance criterion measures how much a risk-averse investor is willing to pay for conditioning on better exchange rate forecasts. We report \mathcal{P} in annualized basis points (*bps*).¹⁷

In the context of mean-variance analysis, perhaps the most commonly used performance measure is the Sharpe ratio (\mathcal{SR}). The realized \mathcal{SR} is equal to the average excess return of a portfolio divided by the standard deviation of the portfolio returns. Note that \mathcal{SR} is a performance measure that is specific to a model, whereas \mathcal{P} is a relative performance measure for a model relative to the benchmark. We assess the statistical significance of economic value using the Ledoit and Wolf (2008) bootstrap two-sided test of whether the Sharpe ratio of a model is different from that of the RW benchmark (see, e.g., Thornton and Valente, 2012).¹⁸ Finally, we also report the maximum drawdown (\mathcal{MDD}), which is the maximum cumulative loss from the strategy's peak to the following trough. A reasonably low \mathcal{MDD} is indicative of the success of an asset allocation strategy because large drawdowns usually lead to fund

in foreign exchange.

¹⁷Note that the Goetzmann, Ingersoll, Spiegel and Welch (2007) performance fee \mathcal{P} defined above gives very similar results to the Fleming, Kirby and Ostdiek (2001) performance fee based on quadratic utility that is often used in the literature.

¹⁸Note that the Ledoit and Wolf (2008) approach does not account for the effects of uncertainty associated with the estimation of forecast model parameters. In contrast, the McCracken and Valente (2012) bootstrap-based approach does account for estimation error in the asymptotic variance of the test statistic. However, the latter approach is not feasible in our context since estimation with the elastic net is computationally intensive and we cannot repeat out-of-sample estimation hundreds or thousands of times as would be required by the McCracken and Valente (2012) test.

redemptions.

4.4 Transaction Costs

A realistic assessment of the profitability of a dynamic trading strategy takes into account the effect of transaction costs. We account for this effect in two ways. First, we compute performance measures for the realized returns net of transaction costs using the bid-ask spread of each spot exchange rate taken from *Datastream*. The net realized return on the investor's portfolio is defined as:

$$r_{p,t+1}^{net} = w_t' r_{t+1} + (1 - w_t') r_f - \tau_{t+1} \quad (19)$$

$$\tau_{t+1} = \sum_{j=1}^N \tau_{j,t+1} (w_{j,t+1} - w_{j,t+1}^-), \quad (20)$$

where τ_{t+1} is the total proportional transaction cost for the portfolio at time $t + 1$, $\tau_{j,t+1}$ is the proportional transaction cost for each asset $j \leq N$ at time $t + 1$, and $w_{j,t+1}^- = w_{j,t} (1 + r_{j,t+1}) / (1 + r_{p,t+1})$.

The proportional transaction cost $\tau_{j,t+1}$ for each asset j is computed as follows. The net return from buying currency j at time t and selling at time $t + 1$ at the spot rate is equal to $s_{j,t+1}^b - s_{j,t}^a = s_{j,t+1} - s_{j,t} - \tau_{j,t+1}$, where $s_{j,t+1}^b$ is the bid-quote for spot rate j at time $t + 1$, $s_{j,t}^a$ is the ask-quote for the spot rate j at time t , $s_{j,t+1}$ and $s_{j,t}$ are the mid-quotes, $\tau_{j,t+1} = \ln \frac{(1+c_{j,t+1})}{(1-c_{j,t+1})}$, and $c_{j,t+1} = 0.5 (S_{j,t+1}^a - S_{j,t+1}^b) / S_{j,t+1}$ is the one-way proportional transaction cost (e.g., Neely, Weller and Ulrich, 2009). Note that upper case $S_{j,t}$ is the spot exchange rate and lower case $s_{j,t}$ is $s_{j,t} = \ln S_{j,t}$.

Similarly, the net return from selling a currency at time t and buying it at time $t + 1$ at the spot rate is equal to $s_{j,t}^b - s_{j,t+1}^a = -(s_{j,t+1} - s_{j,t} - \tau_{j,t+1})$, where $\tau_{j,t+1} = \ln \frac{(1-c_{j,t+1})}{(1+c_{j,t+1})}$.¹⁹ This analysis implies that in Eq. (20) above, if $w_{j,t+1} - w_{j,t+1}^- > 0$ then $\tau_{j,t+1} = \ln \frac{(1+c_{j,t+1})}{(1-c_{j,t+1})}$, and if $w_{j,t+1} - w_{j,t+1}^- < 0$ then $\tau_{j,t+1} = \ln \frac{(1-c_{j,t+1})}{(1+c_{j,t+1})}$. In the tables, we report \mathcal{SR}_τ and \mathcal{P}_τ , which are measures of the Sharpe ratio and the performance fee that directly account for these transaction costs.

Second, we calculate the break-even proportional transaction cost, τ^{be} , that renders investors indifferent between two strategies (e.g., Della Corte, Sarno and Tsiakas, 2009). The

¹⁹The derivation is as follows: $\frac{S_{t+1}^b}{S_t^a} = \frac{S_{t+1} - 0.5(S_{t+1}^a - S_{t+1}^b)}{S_t + 0.5(S_t^a - S_t^b)} = \frac{S_{t+1} \left(1 - \frac{0.5(S_{t+1}^a - S_{t+1}^b)}{S_{t+1}}\right)}{S_t \left(1 + \frac{0.5(S_t^a - S_t^b)}{S_t}\right)} = \frac{S_{t+1}(1-c_{t+1})}{S_t(1+c_t)}$. Then, in logs: $s_{t+1}^b - s_t^a = s_{t+1} - s_t - \ln \frac{(1+c_t)}{(1-c_{t+1})}$. Similarly, $s_t^b - s_{t+1}^a = -(s_{t+1} - s_t - \tau_{t+1})$, where $\tau_{t+1} = \ln \frac{(1-c_t)}{(1+c_{t+1})}$.

proportional transaction cost investors pay at each time $t + 1$ is equal to: $v_{t+1} = \sum_{j=1}^N |w_{j,t+1} - w_{j,t+1}^-|$. Then it is straightforward to show that the τ^{be} that makes investors indifferent between a strategy that delivers portfolio returns $r_{p,t}^*$ and the benchmark strategy that delivers $r_{p,t}^b$ is equal to:

$$\tau^{be} = \frac{\bar{r}_p^* - \bar{r}_p^b}{\bar{v}^* - \bar{v}^b},$$

where \bar{r}_p^* , \bar{r}_p^b , \bar{v}^* and \bar{v}^b are sample means across time. In comparing the E-KS dynamic strategy with the benchmark RW strategy, an investor who pays an actual transaction cost lower than τ^{be} will prefer the E-KS strategy. Since τ^{be} is a proportional cost paid every time the portfolio is rebalanced, we report τ^{be} in monthly basis points. Also note that in the tables we report the τ^{be} value only when the performance fee (\mathcal{P}) is positive.

4.5 Further Asset Allocation Strategies

In addition to the mean-variance asset allocation described above, for robustness we also implement two further asset allocation strategies. The first one is momentum, which for currencies has recently been evaluated by Menkhoff, Sarno, Schmeling and Schrimpf (2012b). We implement the momentum strategy as follows: every month we rank all nine exchange rates according to their excess FX returns. We then go long on the top three with equal weights and short on the bottom three with equal weights. The momentum strategy produces a long exposure to the currencies that are trending higher, and a short exposure to the currencies that are trending lower.²⁰

The second one is the 1/N strategy of DeMiguel, Garlappi and Uppal (2009). This strategy simply sets an equal weight on all (N) assets in a portfolio every month. Hence it requires neither forecasting the asset returns nor optimization in generating the portfolio weights. DeMiguel, Garlappi and Uppal (2009) find that for equities the 1/N strategy outperforms a number of standard asset allocation strategies. Our analysis examines whether this is also the case for exchange rates. Both momentum and the 1/N strategy are straightforward to implement and we use them as further benchmarks in assessing the economic value of the predictive ability of the E-KS model.

²⁰We have also designed a carry trade strategy in similar fashion to the momentum strategy: go long every month on the three currencies with the highest interest rates, and short on the three currencies with the lowest interest rates. However, the performance of this carry trade is practically identical to the mean-variance random walk model. Hence we do not report the results but they are available upon request.

5 Empirical Results

5.1 Data on Exchange Rates and Economic Fundamentals

Our empirical analysis uses spot and forward exchange rates as well as a set of macroeconomic variables for nine exchange rates relative to the US dollar (USD): the Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), Deutsche mark/euro (EUR), British pound (GBP), Japanese yen (JPY), Norwegian krone (NOK), New Zealand dollar (NZD) and Swedish krona (SEK).²¹ All variables are monthly. The data sample ranges from January 1976 to June 2012 for a total of 438 monthly observations.

For exchange rates, we use end-of-month spot and one-month forward rates obtained through the *Download Data Program* of the Board of Governors of the Federal Reserve System. The exchange rate is defined as the US dollar price of a unit of foreign currency so that an increase in the exchange rate implies a depreciation of the US dollar. For interest rates, we use end-of-month Eurodeposit rates from *Datastream*.

We obtain seasonally adjusted data on industrial production, consumer price indices and broad money from the OECD *Main Economic Indicators*. We proxy real output by the industrial production index (IPI) since the IPI is generally available at monthly frequency, whereas GDP data are available quarterly. Note that the IPI of Australia, New Zealand and Switzerland are only available quarterly, hence in these cases we obtain monthly observations via linear interpolation. The price level is measured by the consumer price index (CPI), which is published every month, except for Australia and New Zealand where it is available quarterly. The annual inflation rate is computed as the 12-month log difference of the CPI. Broad money refers to the monetary aggregate M1 that is measured in the national currency, except for the United Kingdom for which we use M0. Finally, with the exception of interest rates, we convert all data by taking logs and multiplying by 100.

Table 1 reports descriptive statistics for the following variables: the monthly percent FX returns, Δs_t ; the difference between domestic and foreign interest rates, $i_t - i_t^*$; the difference in the percent change in price levels, $\Delta (p_t - p_t^*)$; the difference in the percent change in money supply, $\Delta (m_t - m_t^*)$; and the difference in the percent change in real output, $\Delta (y_t - y_t^*)$. For our sample period, the monthly sample means of the FX returns range from -0.104% for SEK to 0.306% for JPY (or from -1.248% to 3.672% per annum). The return standard deviations are similar across all exchange rates at about 3% per month (or about 10% per annum). Most FX returns exhibit negative skewness and higher than normal kurtosis. Finally, the exchange

²¹Before the introduction of the euro in January 1999, we use the US dollar-Deutsche mark exchange rate adjusted by the official conversion rate between the Deutsche mark and the euro.

rate return sample autocorrelations are low and decay rapidly. For the economic fundamentals the notable trends are as follows: $i_t - i_t^*$ are highly persistent with long memory; $\Delta(p_t - p_t^*)$ are always negatively skewed; and $\Delta(m_t - m_t^*)$ and $\Delta(y_t - y_t^*)$ have occasionally very high kurtosis.

In Table 2 we report the cross-correlations among the four predictors. The table shows that although the four predictors capture different aspects of the domestic and foreign economies, their correlations tend to be rather high, positive or negative. The highest correlation is between the PPP predictor and the MF predictor as it ranges from 0.358 to 0.797. It is also worth noting that the Taylor rule tends to be negatively correlated with the other three predictors.

5.2 Statistical Evaluation

We assess the statistical performance of the empirical exchange rate models by reporting out-of-sample tests of predictability against the null of the RW. Our analysis focuses on the R_{oos}^2 statistic of Campbell and Thompson (2008) and Welch and Goyal (2008). Recall that a positive R_{oos}^2 implies that the alternative model has lower MSE and hence higher predictive accuracy than the benchmark RW. The significance of the R_{oos}^2 is assessed using the Clark and West (2006, 2007) one-sided t -statistic and the Giacomini and White (2006) two-sided t -statistic. The OOS monthly forecasts are obtained in two ways: (i) with recursive regressions that generate forecasts for the period of January 1987 to June 2012 by successively re-estimating the model parameters every time a new observation is added to the sample; and (ii) with rolling regressions that use a 10-year window that generates forecasts for the same forecasting period.

The R_{oos}^2 statistics are reported in Table 3. Panel A shows the recursive regression results and Panel B the rolling regression results. In addition, Table 4 presents the Giacomini and Rossi (2009) t -statistic for testing for forecast breakdowns. Our main findings can be summarized in the following six results. First, the E-KS model is the only model that has a positive R_{oos}^2 for every single one of the nine exchange rates and its value revolves around 1%. For example, in the recursive regressions, the R_{oos}^2 is significantly positive for six out of nine exchange rates when using the Clark and West (2006, 2007) one-sided t -statistic and five out of nine exchange rates when using the Giacomini and White (2006) two-sided t -statistic. Therefore, overall the E-KS model consistently outperforms the RW.

Second, in contrast to the E-KS model, the plain KS regression estimated with OLS is by far the worst performing model with a negative R_{oos}^2 in all cases, which revolves around -3%

in the recursive regressions and -7% in the rolling regressions. This is consistent with the results of Welch and Goyal (2008), who find that for equity returns the plain KS regression exhibits very poor OOS performance. Indeed, this result is a principal motivation for using the E-KS model.

Third, the E-KS model considerably outperforms both the ridge regression (RR) and the lasso regression (LR). This clearly indicates that it is not shrinkage per se that is driving the results but rather the particular combination of the RR and LR constraints into the elastic-net estimator used for estimating the E-KS. This motivates the use of elastic-net estimation over simpler shrinkage estimation.

Fourth, the individual empirical models (UIP, PPP, MF, TR) consistently produce negative R_{oos}^2 and thus underperform relative to the E-KS model. Of these models, PPP is the best performing model but only for the recursive regressions. This evidence suggests that none of the individual empirical models can consistently deliver reliable exchange rate forecasts across all exchange rates and types of predictive regressions (i.e., rolling vs. recursive).

Fifth, the combined forecasts produce a negative R_{oos}^2 for about half of the exchange rates. Moreover, when the R_{oos}^2 is positive, it tends to be insignificant. Overall, therefore, the combined forecasts perform significantly worse than the E-KS model. Perhaps the best performing forecast combination is the median, but its performance is overwhelmingly dominated by the E-KS.

Finally, sixth, there is strong evidence that the E-KS forecasts do not suffer from forecast breakdowns over the sample period. The Giacomini and Rossi (2006) t -statistic indicates that for seven out of nine exchange rates in the recursive regressions and eight out of nine in the rolling regressions we cannot reject the null hypothesis of no breakdown in forecasting when using the E-KS model. In these cases, the IS MSE of the E-KS is not significantly higher than the OOS MSE leading to rather stable forecasts. Moreover, it is worth noting that in all cases the E-KS model produces the lowest value for the Giacomini and Rossi (2006) t -statistic and hence the lowest deviation of the OOS MSE from the IS MSE.²² In conclusion, these statistical results indicate that the E-KS model has predictive ability that cannot be matched by any of the empirical models or combinations of models, including the benchmark RW.

²²Note that we implement the Giacomini and Rossi (2009) test for all models except for the combined forecasts. This is because the combined forecasts are constructed using exclusively the OOS forecast of the individual models. Since we do not use IS information to form any forecast combination, it is not meaningful to compare the IS MSE to the OOS MSE of these forecast combinations. Hence for the Giacomini and Rossi (2009) test we focus on individual models.

5.3 Economic Evaluation

We assess the economic value of out-of-sample exchange rate predictability by analyzing the performance of dynamically rebalanced portfolios based on the one-month ahead forecasts generated by recursive or rolling regressions. Our empirical analysis focuses on the Sharpe ratio (\mathcal{SR}), and the Goetzmann, Ingersoll, Spiegel and Welch (2007) performance fee (\mathcal{P}). We assess the statistical significance of economic value using the Ledoit and Wolf (2008) bootstrap two-sided test of whether the Sharpe ratio of a model is different from that of the RW benchmark. In the presence of transaction costs, the Sharpe ratio and the performance fee are denoted by \mathcal{SR}_τ and \mathcal{P}_τ respectively. We also compute the maximum drawdown (MDD) and the break-even transaction cost τ^{be} .

Following Della Corte, Sarno and Tsiakas (2009, 2011) our empirical analysis focuses on the maximum expected return strategy as this is the strategy most often used in active currency management. We set a volatility target of $\sigma_p^* = 10\%$ and a degree of RRA $\gamma = 6$. We have experimented with different σ_p^* and γ values and found that qualitatively they have little effect on the asset allocation results discussed below. Table 5 reports the OOS portfolio performance. Panel A shows the recursive regression results and Panel B the rolling regression results. For the rest of our discussion we focus on the recursive regressions unless otherwise stated.

Overall, we find that there is high economic value associated with the E-KS model. For example, without transaction costs, the RW has an \mathcal{SR} of 0.56, but the E-KS delivers an almost double \mathcal{SR} of 0.94. The difference in the two Sharpe ratios is statistically significant with 95% confidence. Furthermore, the performance fee of the E-KS over the RW is 414 annual basis points (*bps*). In other words, a risk-averse investor on average is willing to pay a fee of more than 4% every year to be able to use the E-KS forecasts rather than the RW. In the FX literature, this is higher economic value than any model based on economic fundamentals (for a review see, for example, Della Corte and Tsiakas, 2012).

With transaction costs, the RW delivers $\mathcal{SR}_\tau = 0.38$, whereas the E-KS provides investors with $\mathcal{SR}_\tau = 0.79$ and $\mathcal{P}_\tau = 465$ *bps*. Even with transaction costs, therefore, the E-KS generates a Sharpe ratio that is almost double that of the RW model. It is important to note that the E-KS is the only model with a \mathcal{SR} that is significantly higher than the RW. No other model outperforms the benchmark RW in terms of delivering a statistically significantly higher \mathcal{SR} .

It is also interesting to note that the performance fee of the E-KS slightly increases when accounting for transaction costs. This result holds only for recursive regressions and is driven

by slightly less volatile weights for the E-KS than the RW.²³ A further piece of empirical evidence supporting the good performance of the E-KS model is the massive value of the break-even transaction cost, which is equal to 1526 *bps* per month. Finally, the E-KS has the lowest MDD of all models at 9.4% (together with the 1/N strategy). In short, therefore, the E-KS model considerably outperforms the RW benchmark net of transaction costs. This indicates that exchange rates are to a large extent predictable and this predictability can lead to high tangible economic gains for investors.

The rest of the empirical models perform much worse than the E-KS and often worse than the RW as well. For example, the plain KS regression estimated with OLS delivers $\mathcal{SR}_\tau = 0.16$ and $\mathcal{P}_\tau = -249$ *bps*. When using ridge regression, the KS delivers $\mathcal{SR}_\tau = 0.38$ and $\mathcal{P}_\tau = 6$ *bps*. When using lasso regression, the KS delivers $\mathcal{SR}_\tau = 0.37$ and $\mathcal{P}_\tau = -6$ *bps*. These results provide ample justification for using the elastic-net estimation method for the kitchen sink regression.

Most of the individual empirical exchange rate models also produce a Sharpe ratio lower than the RW and a negative performance fee. This is consistent with a large literature in international finance that finds that economic fundamentals do not outperform the benchmark random walk. Note that the UIP and PPP perform slightly better than the RW, but only for the recursive regressions. For the rolling regressions, instead, UIP and PPP strongly underperform the RW. Finally, the two alternative asset allocation strategies, momentum and 1/N, also considerably underperform relative to the RW.

Turning to the forecast combinations, most of them deliver an \mathcal{SR}_τ slightly higher than the RW, the highest being 0.50 compared to 0.38 for the RW. However, none of these \mathcal{SR}_τ are significantly higher than the RW when applying the Ledoit and Wolf (2008) test. Moreover, the \mathcal{P}_τ tends to be positive ranging from 66 *bps* to 145 *bps*, where the highest value of the range is for the median combination. The τ^{be} for the forecast combinations are all under 20 *bps* per month. Note that the performance of the combined forecasts is worse for rolling regressions, where, for example, the τ^{be} revolves around 3 *bps*. In short, therefore, the combined forecasts slightly outperform the RW, but not significantly, and their performance is consistently far worse than the E-KS. In light of this evidence, we conclude that the efficient kitchen-sink regression generates forecasts that deliver high economic value to an investor actively managing a currency portfolio. The performance of the E-KS portfolios is significantly better than the

²³For recursive regressions, the effect of transaction costs is to lower the portfolio returns of the E-KS by 1.84% and of the RW by 2.03% (not reported in the tables). Furthermore, the turnover is slightly lower for the E-KS than for the RW. This explains why the performance of the E-KS relative to the RW slightly improves in the presence of transaction costs. As we will see in a following section, the recursive betas of the E-KS are smooth over time, which leads to slightly less volatile weights than the RW. In contrast, for rolling regressions, transaction costs penalize more the E-KS than the RW.

benchmark RW as well as a series of empirical exchange rate models or combinations of models.

To provide a visual illustration of the results, Figures 1 and 2 plot the cumulative wealth of select out-of-sample dynamic investment strategies (solid red line) relative to the RW benchmark (dashed blue line). Figure 1 shows the cumulative wealth of individual empirical exchange rate models and Figure 2 of forecast combinations. Both figures use recursive predictive regressions and portfolio returns without transaction costs. Initial wealth is set at \$1, which grows at the monthly return of the portfolio strategy generated by the out-of-sample forecasts of each model. The figures show emphatically the superior performance of the E-KS regression, which dominates the RW, when all other models fail to do so convincingly. The figures also provide an interesting view as they illustrate at which time period a particular model performs well. This is important because our sample includes the crisis period that started with the credit crunch in July 2007 and culminated with the collapse of Lehman Brothers in September 2008. For example, we can see that the cumulative wealth of the E-KS regression falls during the crisis but after 2009 it increases again as the E-KS continues to outperform the RW.

5.4 Comparing Statistical and Economic Gains

Recall that, among all empirical models, the E-KS model has the highest R_{oos}^2 , which is always positive and predominantly significant but revolves at a value of around 1%, which is seemingly low. Therefore, it is important to note that, as shown by Campbell and Thompson (2008), a low positive R_{oos}^2 can still generate large economic benefits for investors in the context of dynamic trading strategies. The main reason for this is that “...the correct way to judge the magnitude of R^2 is to compare it with the squared Sharpe ratio...”²⁴ because the proportional increase in the expected return is approximately equal to the ratio of the R_{oos}^2 over the squared Sharpe ratio (\mathcal{SR}^2).²⁵ This implies that a low positive R_{oos}^2 for monthly regressions is consistent with a substantially higher monthly expected return.

Consider the following example. The average E-KS R_{oos}^2 across the nine exchange rates is 0.66% for recursive regressions and 1.04% for rolling regressions. At the same time, the E-KS annualized \mathcal{SR} is 0.94 and 0.95 for recursive and rolling regressions respectively. It is then straightforward to show that that a mean-variance investor can increase the average monthly portfolio return by 8.97% (recursive regressions) or 13.83% (rolling regressions).²⁶

²⁴See Campbell and Thompson (2008), page 1525.

²⁵Specifically, when moving from the unconditional forecast of the expected return to a conditional forecast, the proportional increase in the expected return is $\left(\frac{R_{oos}^2}{1-R_{oos}^2}\right)\left(\frac{1+\mathcal{SR}^2}{\mathcal{SR}^2}\right)$, which is approximately equal to $\frac{R_{oos}^2}{\mathcal{SR}^2}$, when R_{oos}^2 and \mathcal{SR}^2 are both small.

²⁶The two squared monthly Sharpe ratios are $(0.94/\sqrt{12})^2 = 0.0736 = 7.36\%$ for recursive regressions and

Furthermore, an advantage of dynamic asset allocation is that it is by design multivariate thus exploiting predictability in all exchange rates rather than focusing on one exchange rate at a time as done in our statistical analysis. Overall, these arguments suggest that modest predictive ability can plausibly generate large economic gains.

5.5 Revealing the Information in Economic Fundamentals

A large body of previous empirical research has documented the difficulty of establishing a relation between economic fundamentals and movements in the exchange rate. One explanation of this difficulty is parameter instability in the predictive regressions, which can be manifested by high variance for the OLS estimator and hence high variation in the period-by-period OOS beta estimates. Parameter instability makes economic fundamentals appear to be disconnected from the exchange rate when in fact they are connected (e.g., Rossi, 2006).

This explanation is consistent with the Bacchetta and van Wincoop (2004, 2013) scapegoat theory of exchange rates in the following way. In the presence of uncertainty about the effect of economic fundamentals on the exchange rate, market participants have a tendency to blame a particular macroeconomic variable (the scapegoat) for the observed exchange rate movements. The blame may quickly shift from one to another variable as new exchange rate observations arrive that seem to be consistent with the second variable. In the scapegoat theory, this may lead to “rational confusion,” which tends to favor the random walk model even if fundamentals are in fact connected to the exchange rate.²⁷

A second explanation is due to the asset pricing model of Engel and West (2005). This model shows that if exchange rates are related to economic fundamentals, they can still appear to follow a random walk if the discount factor is close to one and economic fundamentals are near unit-root processes. Under certain conditions, therefore, exchange rates manifest random walk behavior, but are still consistent with an asset pricing model that links fundamentals to exchange rates.

Our empirical analysis shows that economic fundamentals are directly related to exchange rate movements. The key to revealing the strength of this relation is replacing OLS with the elastic net in estimating a kitchen-sink regression that incorporates several economic fundamentals. The result is a set of parameter estimates (betas), which are more informative than OLS in the sense that they lead to superior forecasts. The elastic net applied to a kitchen-

$(0.95/\sqrt{12})^2 = 0.0752 = 7.52\%$ for rolling regressions. Then the proportional increase in the expected return is $0.66\%/7.36\% = 8.97\%$ and $1.04\%/7.52\% = 13.83\%$ respectively.

²⁷The survey evidence of Cheung and Chinn (2001) based on questionnaires sent to FX traders provides support for the scapegoat theory.

sink regression is thus an effective way of generating economically meaningful forecasts for exchange rates.

This result is illustrated in Figures 3 and 4. Figure 3 plots the out-of-sample variation in the betas of the plain kitchen-sink regression estimated with OLS for each exchange rate. Figure 4 does the same for the efficient kitchen-sink regression estimated with the elastic net. In both cases the betas are shown for the following regressors: UIP in blue, PPP in red, monetary fundamentals in black and the Taylor rule in green. The out-of-sample monthly forecasts are obtained with recursive regressions that generate forecasts for the period of January 1987 to June 2012. Note that, for these figures only, the regressors have been standardized by their sample standard deviation. This is done so that each beta can be interpreted as capturing the effect on the exchange rate return of a one standard deviation movement in the regressors. Standardization makes the betas across different regressors directly comparable.²⁸

The first result to observe is that the scale of the betas for the plain KS is up to 1000 times larger than the E-KS. The betas of the KS model range from about -0.5 to $+1$ (except for JPY for which the scale is even larger). Instead, the betas of the E-KS model range from -0.0015 to 0.0015 . It is clear, therefore, that shrinkage via the elastic net leads to an enormous reduction in the variation of the betas. This is especially evident in the case of monetary fundamentals (black line), which are known to be notoriously unstable in the literature.

Furthermore, it is interesting to note that the betas of the plain KS regression are very unstable at the beginning of the OOS period. However, as more exchange rate observations are used, the betas seem to stabilize and, in fact, converge towards zero. Hence it appears visually as though longer samples lead to both parameter stability and shrinkage. In other words, at the beginning of the OOS period, the OLS parameter estimates exhibit a large divergence from the more stable estimates of the elastic net. However, as we go through the OOS period, the OLS and the elastic-net parameter estimates converge to similar values. Note, however, that the variance of many betas seems to increase around 2008 at the height of the financial crisis.

Finally, the E-KS betas tend to have the sign predicted by theory or well-established empirical deviations from it. Specifically, UIP tends to have a negative sign, as expected by the well-known forward bias; PPP tends to have a positive sign, which implies that currencies are overvalued or undervalued as predicted by PPP; MF tends to have a positive sign, meaning

²⁸Note that the standardization of the regressors is only for illustrating the betas in the figures and has not been used in estimation or forecasting. Therefore, the standardization does not introduce look-ahead bias. Quite simply, we can think of this as dividing each regressor by a constant specific to that regressor and hence scaling the betas by the same constant.

that exchange rate movements on average follow monetary fundamentals; and the Taylor rule betas revolve around zero. In conclusion, the figures illustrate that the E-KS model leads to a drastic reduction in the size and variability of the betas, while maintaining their expected sign. This is a key feature of the E-KS model in efficiently capturing the predictive information of economic fundamentals for exchange rates.

6 Robustness and Extensions

We assess the robustness of exchange rate predictability by extending our analysis in two directions: first, we implement long-horizon predictive regressions and, second, we use quarterly returns. For both robustness checks we focus on the statistical evaluation of the forecasting models using recursive regressions.

6.1 Long-Horizon Predictive Regressions

The long-horizon predictive regressions condition on monthly regressors for forecasting exchange rate returns over horizons of 3-months, 6-months and 12-months ahead. The results are reported in Panels A , B and C, respectively, of Table 6. Our main findings are similar to those for a one-month predictive horizon and can be summarized as follows. The E-KS model is the by far the best performing forecasting model at the 3-month and 6-month horizons. For example, at the 3-month horizon, the E-KS model has a positive R_{oos}^2 for eight out of nine exchange rates, which is significant for all eight exchange rates using the Clark and West (2006, 2007) one-sided t -statistic, while remaining significant for five exchange rates using the Giacomini and White (2006) two-sided t -statistic. However, at the 12-month horizon, it is the combined forecasts that perform the best. For instance, the $MSE(12)$ forecast combination delivers a positive and significant R_{oos}^2 for all exchange rates at the 12-month horizon. Finally, for all horizons the best individual model remains the PPP when using recursive regressions.

6.2 Quarterly Returns

We conclude our empirical analysis by implementing predictive regressions that condition on quarterly regressors for forecasting one-quarter ahead returns. The results shown in Table 7 confirm the good forecasting performance of the E-KS model also for quarterly returns. In particular, the E-KS model delivers a positive R_{oos}^2 for all nine exchange rates, which is significant for six cases using the Clark and West (2006, 2007) one-sided t -statistic, while remaining significant for five cases using the Giacomini and White (2006) two-sided t -statistic.

In short, therefore, our empirical evidence allows us to conclude that the efficient kitchen-sink regression estimated with the elastic net has strong predictive ability for exchange rate returns. This finding tends to be robust across the G10 currencies, different ways of assessing statistical significance, different forecasting windows (recursive vs. rolling), different predictive horizons and different data frequencies.

7 Conclusion

Forming reliable out-of-sample forecasts for exchange rates by conditioning on economic fundamentals has been an important challenge in international finance research for thirty years. In general, there have been three approaches by the literature. One approach is to estimate a particular empirical exchange rate model conditioning on one type of economic fundamental (e.g., interest rates). A second approach is to incorporate a number of economic fundamentals into a single predictive regression, known as the kitchen-sink regression. A third approach involves combining the forecasts of several predictive regressions, each conditioning on one predictor, into a single forecast combination. After three decades of research, however, empirical success remains elusive for all three approaches and, as a result, there is widespread support for the random walk (no predictability) benchmark.

This paper shows that a kitchen-sink regression can indeed lead to reliable forecasts when we implement the elastic-net shrinkage estimator. We term this the efficient kitchen-sink regression because it efficiently utilizes the predictive information from several macroeconomic variables by shrinking all parameter estimates towards zero. The elastic-net estimator combines the benefits of two popular shrinkage estimators: ridge regression and lasso regression. This results in an improvement in the variance-bias tradeoff and leads to superior performance by reducing the effect of less informative predictors in out-of-sample forecasting.

Using a range of statistical and economic measures of predictability, we show that the efficient kitchen-sink regression outperforms a number of alternative models. These include: (i) the random walk benchmark; (ii) four individual exchange rate models based on uncovered interest parity, purchasing power parity, monetary fundamentals and the Taylor rule; (iii) the plain kitchen-sink regression estimated with ordinary least squares; (iv) the kitchen-sink regression estimated with the simpler shrinkage estimators of ridge regression and lasso regression; (v) standard forecast combinations based on the mean squared error of the individual models; and (vi) popular ad-hoc currency strategies such as momentum and the 1/N strategy.

In conclusion, we find that the key to revealing some the predictive information in economic fundamentals is to implement the efficient approach of a kitchen-sink regression estimated with

the elastic-net shrinkage method. Looking through the lens of this methodology, the empirical evidence provides unequivocal support for the view that exchange rates are predictable by conditioning on economic fundamentals. It is likely, therefore, that the lack of empirical success of a large body of previous research is not because economic fundamentals are devoid of predictive information. Rather it may be that a framework based on simple predictive regressions typically conditioning on one fundamental and estimated with ordinary least squares may not be the way to bring to light the predictive ability of economic fundamentals for exchange rates.

Appendix A: Estimation of the Efficient Kitchen-Sink Regression

We estimate the efficient kitchen-sink regression using the elastic-net shrinkage estimator described in Equation (2) by implementing the following procedure. Given s_1 and s_2 , minimizing the penalized least squares is equivalent to minimizing the Lagrangean form:

$$\frac{1}{2} \sum_{t=1}^{T-1} \left(\Delta s_{t+1} - \alpha - \sum_{j=1}^K \beta_j x_{j,t} \right)^2 + \lambda_1 \sum_{j=1}^K |\beta_j| + \lambda_2 \sum_{j=1}^K \beta_j^2, \quad (\text{A.1})$$

where λ_1 and λ_2 are positive constants. In matrix form, Zou and Hastie (2005) show that this is equivalent to minimizing:

$$\beta^T \left(\frac{X^T X + \lambda_2 I}{1 + \lambda_2} \right) \beta - 2 \Delta s^T X \beta + \lambda_1 \sum_{j=1}^K |\beta_j|, \quad (\text{A.2})$$

where β is the $K \times 1$ vector of regression coefficients, Δs is the $T \times 1$ vector of dependent variables, and X is the $T \times K$ matrix of independent variables.²⁹

We estimate the elastic-net regression with the coordinate descent algorithm of Friedman, Hastie, Hofling and Tibshirani (2007). This algorithm is based on the shooting strategy of Fu (1998) and is similar to the LARS algorithm of Efron, Hastie, Johnstone and Tibshirani (2004). The coordinate descent algorithm involves two stages. In the first stage, we estimate the two tuning parameters λ_1 and λ_2 (or equivalently s_1 and s_2) with five-fold cross-validation. Cross-validation is a data-driven method for determining tuning parameters, which has been extensively used in statistics (e.g., Arlot and Celisse, 2010) and finance (e.g., DeMiguel, Garrappi, Nogales and Uppal, 2009). In the second stage, given the estimates of the tuning parameters $(\hat{\lambda}_1, \hat{\lambda}_2)$, we estimate the vector of regression coefficients β . The two-stage estimation involves the following steps:

1. Given an initial set of values for λ_1 and λ_2 , say $(\lambda_1^{(1)}, \lambda_2^{(1)})$, we randomly partition the in-sample data into five subsamples of roughly equal size, each with M_d observations for $d \leq 5$.
2. For each subsample d , we estimate the regression coefficients by minimizing Eq. (A.2) with $\lambda_1 = \lambda_1^{(1)}$ and $\lambda_2 = \lambda_2^{(1)}$ using all the in-sample data observations, which are *not*

²⁹Equation (A.2) represents the elastic net. If instead we replace the term $\frac{X^T X + \lambda_2 I}{1 + \lambda_2}$ by $X^T X$, then the equation would represent the lasso regression. Furthermore, note that $\frac{X^T X + \lambda_2 I}{1 + \lambda_2} = \delta X^T X + (1 - \delta)I$ is a shrinkage estimator of the covariance matrix with $\delta = 1/(1 + \lambda_2)$ (see, e.g., Ledoit and Wolf, 2004).

included in subsample d . Estimation in this step is based on the coordinate descent algorithm of Friedman, Hastie, Hofling and Tibshirani (2007).

3. We use the estimated regression coefficients $\hat{\beta}^{(d)}$ to generate in-sample forecasts of the FX returns for the particular subsample d .
4. We repeat this procedure for the remaining subsamples and calculate the mean squared error (MSE) of the forecasts, which is conditional on $(\lambda_1^{(1)}, \lambda_2^{(1)})$.
5. We repeat steps 2-4 for different values of λ_1 and λ_2 , say $(\lambda_1^{(2)}, \lambda_2^{(2)})$, $(\lambda_1^{(3)}, \lambda_2^{(3)})$, \dots . Then, we select the combination of tuning parameters $(\hat{\lambda}_1, \hat{\lambda}_2)$ that gives the lowest MSE.
6. Finally, given the tuning parameters $(\hat{\lambda}_1, \hat{\lambda}_2)$, we implement the coordinate descent algorithm of Friedman, Hastie, Hofling and Tibshirani (2007) to estimate $\hat{\beta}$.

Appendix B: Examples of the The Engel and West (2005) Asset Pricing Model

The Engel and West (2005) asset pricing model nests the empirical models on which the predictors in the kitchen-sink regression are based. In this Appendix, we discuss two examples in detail.

B.1 Monetary Fundamentals

The first model to be nested by the Engel and West (2005) present value relation is the monetary fundamentals model, which assumes that the money market relation is described by:

$$m_t = p_t + \gamma y_t - \alpha i_t + v_{m,t}, \quad (\text{B.1})$$

where m_t is the log of the domestic money supply, p_t is the log of the domestic price level, $\gamma > 0$ is the income elasticity of money demand, y_t is the log of the domestic national income, $\alpha > 0$ is the interest rate semi-elasticity of money demand, i_t is the domestic nominal interest rate and $v_{m,t}$ is a shock to domestic money demand. A similar equation holds for the foreign economy, where the corresponding variables are denoted by m_t^* , p_t^* , y_t^* , i_t^* and $v_{m,t}^*$. We assume that the parameters $\{\gamma, \alpha\}$ of the foreign money demand are the same as the domestic parameters.

The nominal exchange rate is equal to its purchasing power parity (PPP) value plus the real exchange rate q_t :

$$s_t = p_t - p_t^* + q_t. \quad (\text{B.2})$$

The interest parity condition is given by:

$$E_t s_{t+1} - s_t = i_t - i_t^* + \rho_t, \quad (\text{B.3})$$

where ρ_t is the deviation from the uncovered interest parity (UIP) condition that is based on rational expectations and risk neutrality. Hence ρ_t can be interpreted either as an expectational error or a risk premium.

Using Equations (B.1) to (B.3) for the domestic and foreign economies and re-arranging, we get:

$$s_t = \frac{1}{1 + \alpha} [m_t - m_t^* - \gamma (y_t - y_t^*) + q_t - (v_{m,t} - v_{m,t}^*) - \alpha \rho_t] + \frac{\alpha}{1 + \alpha} E_t s_{t+1}. \quad (\text{B.4})$$

This equation takes the form of the original model in Equation (7), where the discount factor

is given by $b = \alpha/1 + \alpha$, the observable fundamentals are $f_{1,t} = m_t - m_t^* - \gamma(y_t - y_t^*)$, and the unobservable fundamentals are $z_{1,t} = q_t - (v_{m,t} - v_{m,t}^*)$ and $z_{2,t} = -\rho_t$.

B.2 Taylor Rule

The second model to be nested by the Engel and West (2005) present value relation is the Taylor (1993) rule, where the home country is assumed to set the short-term nominal interest rate according to:

$$i_t = \bar{i} + \beta_1 y_t^g + \beta_2 (\pi_t - \bar{\pi}) + v_t, \quad (\text{B.5})$$

where \bar{i} is the target short-term interest rate, y_t^g is the output gap measured as the percent deviation of real GDP from an estimate of its potential level, π_t is the inflation rate, $\bar{\pi}$ is the target inflation rate and v_t is a shock. The Taylor rule postulates that the central bank raises the short-term nominal interest rate when output is above potential output or inflation rises above its desired level.

The foreign country is assumed to follow a Taylor rule that explicitly targets exchange rates (e.g., Clarida, Gali and Gertler, 1998):

$$i_t^* = -\beta_0 (s_t - \bar{s}_t) + \bar{i} + \beta_1 y_t^{*g} + \beta_2 (\pi_t^* - \bar{\pi}) + v_t^*, \quad (\text{B.6})$$

where $0 < \beta_0 < 1$ and \bar{s}_t is the target exchange rate. For simplicity, we assume that the home and foreign countries target the same interest rate, \bar{i} , and the same inflation rate, $\bar{\pi}$. The rule indicates that the foreign country raises interest rates when its currency depreciates relative to the target.³⁰

We assume that the foreign central bank targets the PPP level of the exchange rate:

$$\bar{s}_t = p_t - p_t^*. \quad (\text{B.7})$$

Taking the difference between the home and foreign Taylor rules, using interest parity (B.3), substituting the target exchange rate and solving for s_t gives:

$$s_t = \frac{\beta_0}{1 + \beta_0} (p_t - p_t^*) - \frac{1}{1 + \beta_0} [\beta_1 (y_t^g - y_t^{*g}) + \beta_2 (\pi_t - \pi_t^*) + v_t - v_t^* + \rho_t] + \frac{1}{1 + \beta_0} E_t s_{t+1}. \quad (\text{B.8})$$

This equation also has the general form of the present value model in Equation (7), where the discount factor is $b = 1/1 + \beta_0$, $f_{1,t} = p_t - p_t^*$ and $z_{2,t} = -[\beta_1 (y_t^g - y_t^{*g}) + \beta_2 (\pi_t - \pi_t^*) + v_t - v_t^* + \rho_t]$.

³⁰The argument still follows if the home country also targets exchange rates. It is standard to omit the exchange rate target from Equation (B.5) on the interpretation that US monetary policy has essentially ignored exchange rates (see, Engel and West, 2005).

References

- Akram, F., D. Rime, and L. Sarno (2008). “Arbitrage in the Foreign Exchange Market: Turning on the Microscope,” *Journal of International Economics* **76**, 237–253.
- Arlot, S., and A. Celisse (2010). “A Survey of Cross-Validation Procedures for Model Selection,” *Statistics Surveys* **4**, 40–79.
- Bacchetta, P., and E. van Wincoop (2004). “A Scapegoat Model of Exchange-Rate Fluctuations,” *American Economic Review Papers and Proceedings* **94**, 114–118.
- Bacchetta, P., and E. van Wincoop (2013). “On the Unstable Relationship Between Exchange Rates and Macroeconomic Fundamentals,” *Journal of International Economics* **91**, 18–26.
- Bai, J., and S. Ng (2008). “Forecasting Economic Time Series using Targeted Predictors,” *Journal of Econometrics* **146**, 304–317.
- Bates, J.M., and C.W.J. Granger (1969). “The Combination of Forecasts,” *Operations Research Quarterly* **20**, 451–468.
- Bilson, J.F.O. (1981). “The ‘Speculative Efficiency’ Hypothesis,” *Journal of Business* **54**, 435–451.
- Burnside, C., M. Eichenbaum, I. Kleshchelski, and S. Rebelo (2011). “Do Peso Problems Explain the Returns to the Carry Trade?” *Review of Financial Studies* **24**, 853–891.
- Campbell, J.Y., and R.J. Shiller (1987). “Cointegration and Tests of Present Value Models,” *Journal of Political Economy* **95**, 1062–1088.
- Campbell, J.Y., and R.J. Shiller (1988). “Stock Prices, Earnings, and Expected Dividends,” *Journal of Finance* **43**, 661–676.
- Campbell, J.Y., and S.B. Thompson (2008). “Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?” *Review of Financial Studies* **21**, 1509–1531.
- Cheung, Y.-W., and M.D. Chinn (2001). “Currency Traders and Exchange Rate Dynamics: A Survey of the US Market,” *Journal of International Money and Finance* **20**, 439–471.
- Clarida, R., J. Gali, and M. Gertler (1998). “Monetary Policy Rules in Practice. Some International Evidence,” *European Economic Review* **42**, 1033–1067.
- Clark, T.E., and M.W. McCracken (2013). “Advances in Forecast Evaluation,” in Elliott, G., and A. Timmermann (eds.), *Handbook of Economic Forecasting*, Volume 2, Amsterdam: North Holland.
- Clark, T.E., and K.D. West (2006). “Using Out-of-Sample Mean Squared Prediction Errors to Test the Martingale Difference Hypothesis,” *Journal of Econometrics* **135**, 155–186.
- Clark, T.E., and K.D. West (2007). “Approximately Normal Tests for Equal Predictive Accuracy,” *Journal of Econometrics* **138**, 291–311.
- De Mol, C., D. Giannone, and L. Reichlin (2008). “Forecasting Using a Large Number of Predictors: Is Bayesian Shrinkage a Valid Alternative to Principal Components?” *Journal of Econometrics* **146**, 318–328.

- Della Corte, P., L. Sarno, and I. Tsiakas (2009). “An Economic Evaluation of Empirical Exchange Rate Models,” *Review of Financial Studies* **22**, 3491–3530.
- Della Corte, P., L. Sarno, and I. Tsiakas (2011). “Spot and Forward Volatility in Foreign Exchange,” *Journal of Financial Economics* **100**, 496–513.
- Della Corte, P., L. Sarno, and I. Tsiakas (2012). “Volatility and Correlation Timing in Active Currency Management,” in James, J., L. Sarno and I.W. Marsh (eds.) *Handbook of Exchange Rates*. Hoboken, NJ: Wiley.
- Della Corte, P., and I. Tsiakas (2012). “Statistical and Economic Methods for Evaluating Exchange Rate Predictability,” in James, J., L. Sarno and I.W. Marsh (eds.) *Handbook of Exchange Rates*. Hoboken, NJ: Wiley.
- DeMiguel, V., L. Garlappi, F.J. Nogales, and R. Uppal (2009). “A Generalized Approach to Portfolio Optimization: Improving Performance by Constraining Portfolio Norms,” *Management Science* **55**, 798–812.
- DeMiguel, V., L. Garlappi, and R. Uppal (2009). “Optimal versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy?” *Review of Financial Studies* **22**, 1915–1953.
- Eickmeier, S., and T. Ng (2011). “Forecasting National Activity using Lots of International Predictors: An Application to New Zealand,” *International Journal of Forecasting* **27**, 496–511.
- Efron, B., T. Hastie, I. Johnstone, and R. Tibshirani (2004). “Least Angle Regression,” *The Annals of Statistics* **32**, 407–499.
- Engel, C., N.C. Mark, and K.D. West (2007). “Exchange Rate Models are not as Bad as you Think,” *NBER Macroeconomics Annual 2007*, 381–441.
- Engel, C., and K.D. West (2005). “Exchange Rates and Fundamentals,” *Journal of Political Economy* **113**, 485–517.
- Fama, E.F. (1984). “Forward and Spot Exchange Rates,” *Journal of Monetary Economics* **14**, 319–338.
- Fleming, J., C. Kirby, and B. Ostdiek (2001). “The Economic Value of Volatility Timing,” *Journal of Finance* **56**, 329–352.
- Friedman, J., T. Hastie, H. Hofling, and R. Tibshirani (2007). “Pathwise Coordinate Optimization,” *Annals of Applied Statistics* **1**, 302–332.
- Fu, W.J. (1998). “Penalized Regressions: the Bridge versus the Lasso,” *Journal of Computational and Graphical Statistics* **7**, 397–416.
- Giacomini, R., and B. Rossi (2009). “Detecting and Predicting Forecast Breakdowns,” *Review of Economic Studies* **76**, 669–705.
- Giacomini, R., and H. White (2006). “Tests of Conditional Predictive Ability,” *Econometrica* **74**, 1545–1578.
- Goetzmann, W., J. Ingersoll, M. Spiegel, and I. Welch (2007). “Portfolio Performance Manipulation and Manipulation-Proof Performance Measures,” *Review of Financial Studies* **20**, 1503–1546.

- Groen, J.J.J. (2000). “The Monetary Exchange Rate Model as a Long-Run Phenomenon,” *Journal of International Economics* **52**, 299–319.
- Hodrick, R.J., and E.C. Prescott (1997). “Postwar U.S. Business Cycles: An Empirical Investigation,” *Journal of Money, Credit and Banking* **29**, 1–16.
- Hoerl, A.E., and R.W. Kennard (1970). “Ridge Regression: Biased Estimation for Nonorthogonal Problems,” *Technometrics* **12**, 55–67.
- Jagannathan, R., and T. Ma (2003). “Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps,” *Journal of Finance* **58**, 1651–1684.
- Jorion, P. (1985). “International Portfolio Diversification with Estimation Risk,” *Journal of Business* **58**, 259–278.
- Jorion, P. (1986). “Bayes-Stein Estimation for Portfolio Analysis,” *Journal of Financial and Quantitative Analysis* **21**, 279–292.
- Jorion, P. (1991). “Bayesian and CAPM Estimators of the Means: Implications for Portfolio Selection,” *Journal of Banking Finance* **15**, 717–727.
- Kim, H.H., and N.R. Swanson (2014). “Forecasting Financial and Macroeconomic Variables Using Data Reduction Methods: New Empirical Evidence,” *Journal of Econometrics* **178**, 352–367.
- Korobilis, D. (2013). “Hierarchical Shrinkage Priors for Dynamic Regressions with Many Predictors,” *International Journal of Forecasting* **29**, 43–59.
- Ledoit, O., and M. Wolf (2003). “Improved Estimation of the Covariance Matrix of Stock Returns with an Application to Portfolio Selection,” *Journal of Empirical Finance* **10**, 603–621.
- Ledoit, O., and M. Wolf (2004). “A Well-Conditioned Estimator for Large-Dimensional Covariance Matrices,” *Journal of Multivariate Analysis* **88**, 365–411.
- Ledoit, O., and M. Wolf (2008). “Robust Performance Hypothesis Testing with the Sharpe Ratio,” *Journal of Empirical Finance* **15**, 850–859.
- Lustig, H., N. Roussanov, and A. Verdelhan (2011). “Common Risk Factors in Currency Markets,” *Review of Financial Studies* **24**, 3731–3777.
- Mancini-Griffoli, T., and A. Ranaldo (2011). “Limits to Arbitrage During the Crisis: Funding Liquidity Constraints and Covered Interest Parity,” Swiss National Bank working paper 2010-14.
- Mark, N.C. (1995). “Exchange Rates and Fundamentals: Evidence on Long-Horizon Predictability,” *American Economic Review* **85**, 201–218.
- Mark, N.C. (2009). “Changing Monetary Policy Rules, Learning, and Real Exchange Rate Dynamics,” *Journal of Money, Credit and Banking* **41**, 1047–1070.
- Mark, N.C., and D. Sul (2001). “Nominal Exchange Rates and Monetary Fundamentals: Evidence from a Small Post-Bretton Woods Panel,” *Journal of International Economics* **53**, 29–52.

- McCracken, M.W. (2007). “Asymptotics for Out of Sample Tests of Granger Causality,” *Journal of Econometrics* **140**, 719–752.
- McCracken, M.W., and G. Valente (2012). “Testing the Economic Value of Asset Return Predictability,” Federal Reserve Bank of St. Louis Working Paper 2012-049A.
- Meese, R.A., and K. Rogoff (1983). “Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?” *Journal of International Economics* **14**, 3–24.
- Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2012a). “Carry Trades and Global Foreign Exchange Volatility,” *Journal of Finance* **67**, 681–718.
- Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2012b). “Currency Momentum Strategies,” *Journal of Financial Economics* **106**, 660–684.
- Molodtsova, T., and D.H. Papell (2009). “Out-of-Sample Exchange Rate Predictability with Taylor Rule Fundamentals,” *Journal of International Economics* **77**, 167–180.
- Neely, C.J., P.A. Weller, and J.M. Ulrich (2009). “The Adaptive Markets Hypothesis: Evidence from the Foreign Exchange Market,” *Journal of Financial and Quantitative Analysis* **44**, 467–488.
- Rapach, D.E., J.K. Strauss, and G. Zhou (2010). “Out-of-Sample Equity Premium Prediction: Combination Forecasts and Links to the Real Economy,” *Review of Financial Studies* **23**, 821–862.
- Rapach, D.E., J.K. Strauss, and G. Zhou (2013). “International Stock Return Predictability: What is the Role of the United States?” *Journal of Finance* **68**, 1633–1662.
- Rapach, D.E., and M.E. Wohar (2002). “Testing the Monetary Model of Exchange Rate Determination: New Evidence from a Century of Data,” *Journal of International Economics* **58**, 359–385.
- Rossi, B. (2006). “Are Exchange Rates Really Random Walks? Some Evidence Robust to Parameter Instability,” *Macroeconomic Dynamics* **10**, 20–38.
- Rossi, B. (2013). “Exchange Rate Predictability,” *Journal of Economic Literature* **51**, 1063–1119.
- Stock, J.H., and M.W. Watson (2004). “Combination Forecasts of Output Growth in a Seven-Country Data Set,” *Journal of Forecasting* **23**, 405–430.
- Stock, J.H., and M.W. Watson (2012). “Generalized Shrinkage Methods for Forecasting Using Many Predictors,” *Journal of Business and Economic Statistics* **30**, 481–493.
- Taylor, A.M., and M.P. Taylor (2004). “The Purchasing Power Parity Debate,” *Journal of Economic Perspectives* **18**, 135–158.
- Taylor, J.B. (1993). “Discretion versus Policy Rules in Practice,” *Carnegie-Rochester Conference Series on Public Policy* **39**, 195–214.
- Thornton, D.L., and G. Valente (2012). “Out-of-Sample Predictions of Bond Excess Returns and Forward Rates: An Asset Allocation Perspective,” *Review of Financial Studies* **25**, 3141–3168.

- Tibshirani, R. (1996). “Regression Shrinkage and Selection via the Lasso,” *Journal of the Royal Statistical Society. Series B (Methodological)* **58**, 267–288.
- Timmermann, A. (2006). “Forecast Combinations,” in Elliott, G., C.W.J. Granger, and A. Timmermann (eds.), *Handbook of Economic Forecasting*, Volume 1, Amsterdam: North-Holland.
- Tu, J., and G. Zhou (2011). “Markowitz Meets Talmud: A Combination of Sophisticated and Naive Diversification Strategies,” *Journal of Financial Economics* **99**, 204–215.
- Welch, I., and A. Goyal (2008). “A Comprehensive Look at the Empirical Performance of Equity Prediction,” *Review of Financial Studies* **21**, 1455–1508.
- West, K.D. (1988). “Dividend Innovations and Stock Price Volatility,” *Econometrica* **56**, 37–61.
- Wright, J.H. (2008). “Bayesian Model Averaging and Exchange Rate Forecasting,” *Journal of Econometrics* **146**, 329–341.
- Zou, H., and T. Hastie (2005). “Regularization and Variable Selection via the Elastic Net,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **67**, 301–320.

Table 1. Descriptive Statistics

The table presents descriptive statistics for monthly log exchange rate returns and a set of monthly economic fundamentals. The exchange rate is defined as the US dollar price of a unit of foreign currency so that an increase in the exchange rate implies a depreciation of the US dollar. Δs is the % change in the US dollar exchange rate against the Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), German mark/euro (EUR), British pound (GBP), Japanese yen (JPY), Norwegian krone (NOK), New Zealand dollar (NZD) and Swedish krona (SEK); i is the one-month interest rate; Δp is the % change in the price level; Δm is the % change in the money supply; Δy is the % change in real output; and the asterisk denotes a non-US value. ρ_l is the autocorrelation coefficient for l lags. The data range from January 1976 to June 2012.

		<i>Mean</i>	<i>Std</i>	<i>Skew</i>	<i>Kurt</i>	ρ_1	ρ_3	ρ_6	ρ_{12}
<i>AUD</i>	Δs	-0.047	3.315	-1.221	8.628	0.020	0.062	0.030	-0.070
	$i - i^*$	-0.186	0.285	0.188	4.621	0.957	0.875	0.798	0.645
	$\Delta(p - p^*)$	-0.080	0.392	-0.801	4.519	0.559	0.104	0.151	0.234
	$\Delta(m - m^*)$	-0.291	1.987	0.805	10.721	-0.013	0.110	-0.086	0.342
	$\Delta(y - y^*)$	0.031	0.781	-0.246	3.445	0.302	0.061	0.070	0.008
<i>CAD</i>	Δs	0.005	1.941	-0.646	9.907	-0.060	0.034	-0.104	0.037
	$i - i^*$	-0.064	0.136	-0.222	3.858	0.881	0.720	0.630	0.415
	$\Delta(p - p^*)$	0.006	0.350	-0.715	6.769	0.016	0.085	0.075	0.236
	$\Delta(m - m^*)$	-0.218	1.221	0.222	3.692	-0.123	0.297	0.146	0.321
	$\Delta(y - y^*)$	-0.012	1.068	-0.918	11.383	-0.213	0.067	-0.029	-0.075
<i>CHF</i>	Δs	0.207	3.602	-0.114	4.142	0.020	0.035	-0.064	-0.042
	$i - i^*$	0.242	0.277	0.540	3.795	0.971	0.892	0.836	0.749
	$\Delta(p - p^*)$	0.162	0.413	-0.088	3.801	0.229	-0.061	0.339	0.506
	$\Delta(m - m^*)$	-0.013	2.137	-0.169	4.993	-0.106	0.318	0.345	0.371
	$\Delta(y - y^*)$	0.011	0.926	-0.020	4.307	0.323	-0.062	-0.012	-0.017
<i>EUR</i>	Δs	0.118	3.249	-0.160	3.620	0.004	0.046	-0.030	-0.021
	$i - i^*$	0.107	0.235	-0.146	3.602	0.961	0.877	0.829	0.723
	$\Delta(p - p^*)$	0.127	0.409	-0.746	5.011	0.145	0.017	0.061	0.513
	$\Delta(m - m^*)$	-0.141	2.450	-0.091	5.932	-0.265	0.245	0.257	0.492
	$\Delta(y - y^*)$	0.048	1.806	0.036	8.864	-0.277	0.087	-0.011	-0.007
<i>GBP</i>	Δs	-0.059	3.043	-0.205	4.726	0.077	0.040	-0.048	0.010
	$i - i^*$	-0.168	0.216	-0.767	5.550	0.920	0.765	0.572	0.277
	$\Delta(p - p^*)$	-0.064	0.545	-1.343	8.801	0.089	-0.020	0.238	0.533
	$\Delta(m - m^*)$	-0.049	1.293	0.312	4.793	-0.125	0.204	0.016	0.446
	$\Delta(y - y^*)$	0.132	1.315	0.281	6.565	-0.245	0.058	-0.079	-0.061
<i>JPY</i>	Δs	0.306	3.313	0.341	4.373	0.038	0.038	-0.106	0.029
	$i - i^*$	0.233	0.222	0.419	3.369	0.952	0.832	0.679	0.463
	$\Delta(p - p^*)$	0.203	0.512	-0.453	5.410	0.078	-0.112	0.106	0.414
	$\Delta(m - m^*)$	-0.070	3.081	-0.859	9.073	-0.304	0.393	0.356	0.464
	$\Delta(y - y^*)$	0.051	1.857	2.450	25.116	-0.107	0.027	0.046	-0.067
<i>NOK</i>	Δs	-0.016	3.049	-0.442	4.418	0.022	-0.003	-0.005	-0.060
	$i - i^*$	-0.180	0.269	0.004	3.223	0.934	0.830	0.741	0.488
	$\Delta(p - p^*)$	-0.028	0.533	-0.879	5.437	0.161	0.122	0.135	0.431
	$\Delta(m - m^*)$	-0.374	3.193	-1.429	13.017	-0.305	-0.131	0.176	0.120
	$\Delta(y - y^*)$	-0.022	4.303	0.661	28.126	-0.416	-0.038	-0.006	-0.097
<i>NZD</i>	Δs	-0.060	3.436	-1.160	10.743	0.013	0.195	0.065	-0.080
	$i - i^*$	-0.291	0.341	-1.746	8.092	0.957	0.859	0.736	0.546
	$\Delta(p - p^*)$	-0.158	0.491	-1.574	8.201	0.684	0.338	0.297	0.279
	$\Delta(m - m^*)$	-0.227	3.414	-0.047	3.519	-0.356	0.036	0.213	0.472
	$\Delta(y - y^*)$	0.034	1.117	-1.197	7.639	0.527	0.151	0.043	0.027
<i>SEK</i>	Δs	-0.104	3.261	-0.760	5.998	0.073	0.073	-0.073	-0.024
	$i - i^*$	-0.147	0.270	-0.858	4.383	0.936	0.837	0.748	0.527
	$\Delta(p - p^*)$	-0.039	0.542	-1.076	6.312	0.188	0.055	0.095	0.348
	$\Delta(m - m^*)$	-0.340	3.412	-6.901	100.937	0.052	0.027	-0.020	0.053
	$\Delta(y - y^*)$	0.093	2.788	0.084	37.116	-0.297	0.039	0.039	-0.022

Table 2. Cross-Correlations Among Predictors

The table reports the cross-correlations among the four monthly predictors used in the kitchen-sink regression. These include: (i) a predictor based on uncovered interest parity, $x_{1,t} = f_t - s_t$, where f_t is the one-month forward exchange rate and s_t the spot exchange rate; (ii) a predictor based on purchasing power parity, $x_{2,t} = p_t - p_t^* - s_t$, where p_t is the domestic price level and p_t^* the foreign price level; (iii) a predictor based on monetary fundamentals, $x_{3,t} = (m_t - m_t^*) - (y_t - y_t^*) - s_t$, where m_t is the domestic money supply, m_t^* the foreign money supply, y_t the domestic output and y_t^* the foreign output; and (iv) a predictor based on the Taylor rule, $x_{4,t} = 1.5(\pi_t - \pi_t^*) + 0.1(y_t^g - y_t^{*g}) + 0.1(s_t + p_t^* - p_t)$, where π_t is the domestic inflation, π_t^* the foreign inflation, y_t^g the domestic output gap and y_t^{*g} the foreign output gap. The data range from January 1976 to June 2012.

		x_1	x_2	x_3	x_4
<i>AUD</i>	x_1	1.000			
	x_2	-0.075	1.000		
	x_3	-0.404	0.738	1.000	
	x_4	-0.031	-0.298	-0.179	1.000
<i>CAD</i>	x_1	1.000			
	x_2	0.298	1.000		
	x_3	-0.131	0.527	1.000	
	x_4	-0.155	-0.255	-0.213	1.000
<i>CHF</i>	x_1	1.000			
	x_2	0.280	1.000		
	x_3	0.262	0.662	1.000	
	x_4	-0.157	-0.217	-0.114	1.000
<i>EUR</i>	x_1	1.000			
	x_2	0.160	1.000		
	x_3	0.431	0.456	1.000	
	x_4	-0.184	-0.124	0.003	1.000
<i>GBP</i>	x_1	1.000			
	x_2	0.118	1.000		
	x_3	-0.304	0.474	1.000	
	x_4	0.153	-0.136	-0.091	1.000
<i>JPY</i>	x_1	1.000			
	x_2	0.169	1.000		
	x_3	0.134	0.412	1.000	
	x_4	-0.159	0.026	0.019	1.000
<i>NOK</i>	x_1	1.000			
	x_2	-0.031	1.000		
	x_3	-0.277	0.619	1.000	
	x_4	0.071	-0.089	-0.168	1.000
<i>NZD</i>	x_1	1.000			
	x_2	-0.057	1.000		
	x_3	-0.364	0.797	1.000	
	x_4	-0.133	-0.200	-0.069	1.000
<i>SEK</i>	x_1	1.000			
	x_2	0.313	1.000		
	x_3	-0.330	0.358	1.000	
	x_4	-0.100	-0.048	-0.092	1.000

Table 3. Statistical Evaluation of Exchange Rate Predictability

The table displays the out-of-sample R_{oos}^2 test statistic in percent for a set of empirical exchange rate models and combined forecasts against the null of a random walk (RW). The out-of-sample monthly forecasts are obtained in two ways: with recursive regressions that generate forecasts for the period of January 1987 to June 2012 by successively re-estimating the model parameters every time a new observation is added to the sample (Panel A); and with rolling regressions using a rolling window of 10 years that generate forecasts for the same out-of-sample period (Panel B). The combined forecasts use combinations of five models: the random walk, uncovered interest parity, purchasing power parity, monetary fundamentals and the Taylor rule. a , b , and c denote statistical significance at the 10%, 5%, and 1% level, respectively. The superscripts a , b , and c use the Clark and West (2006, 2007) one-sided t -statistic, whereas the subscripts a , b , and c use the Giacomini and White (2006) two-sided t -statistic.

	Panel A: R_{oos}^2 (%) – Recursive Regressions								
	<i>AUD</i>	<i>CAD</i>	<i>CHF</i>	<i>EUR</i>	<i>GBP</i>	<i>JPY</i>	<i>NOK</i>	<i>NZD</i>	<i>SEK</i>
<i>Efficient Kitchen-Sink Regression</i>	0.986 ^b	0.787 ^c	0.496	0.486	0.504 ^c	0.229	0.598 ^b	1.164 ^b	0.650 ^a
<i>Ridge Regression</i>	0.015	0.023	0.013	0.002	0.224 ^a	0.012	−0.006	0.034	−0.033
<i>Lasso Regression</i>	0.000	−0.028	0.000	0.000	0.149 ^a	−0.072	0.002	0.000	0.000
<i>Plain Kitchen-Sink Regression</i>	−1.781	−0.553	−3.517	−4.533	−3.396	−2.223	−4.456	−2.895	−2.235
<i>Uncovered Interest Parity</i>	0.037	−0.194	−2.813	−2.253	−3.346	−0.195	−0.488	0.172	−0.834
<i>Purchasing Power Parity</i>	−0.280	−0.644	0.983 ^a	0.850 ^b	1.662 ^b	0.670 ^a	0.610 ^a	−0.610	0.995 ^b
<i>Monetary Fundamentals</i>	−0.214	−1.386	−0.721	−0.516	−0.014	−1.340	−0.796	−0.099	−0.126
<i>Taylor Rule</i>	−0.510	0.268	−1.173	−2.990	−2.153	−1.107	−3.471	−0.813	−0.791
<i>Mean Combination</i>	−0.004	0.227	0.549 ^a	−0.160	0.221	0.733 ^a	−0.409	−0.064	0.044
<i>Median Combination</i>	0.176	0.557 ^a	0.110	0.158	−0.282	1.054 ^c	0.048	−0.179	0.064
<i>Trimmed Mean Combination</i>	0.116	0.270	0.293	0.088	0.116	0.846 ^c	0.027	−0.053	0.025
<i>DMSE Combination</i> ($\varphi = 0.90$)	−0.004	0.333	0.354	−0.169	0.354	0.698 ^a	−0.341	−0.034	0.080
<i>DMSE Combination</i> ($\varphi = 0.95$)	−0.004	0.275	0.412	−0.158	0.383	0.663 ^a	−0.359	−0.036	0.073
<i>DMSE Combination</i> ($\varphi = 1.00$)	−0.007	0.233	0.474 ^a	−0.159	0.371	0.630 ^a	−0.376	−0.046	0.051
<i>MSE Combination</i> ($\kappa = 60$)	−0.007	0.232	0.475 ^a	−0.159	0.369	0.611 ^a	−0.370	−0.043	0.050
<i>MSE Combination</i> ($\kappa = 36$)	−0.006	0.224	0.462 ^a	−0.172	0.366	0.596 ^a	−0.374	−0.039	0.045
<i>MSE Combination</i> ($\kappa = 12$)	−0.002	0.188	0.513 ^a	−0.134	0.277	0.560	−0.354	−0.028	0.055

Panel B: R_{oos}^2 (%) – Rolling Regressions									
	<i>AUD</i>	<i>CAD</i>	<i>CHF</i>	<i>EUR</i>	<i>GBP</i>	<i>JPY</i>	<i>NOK</i>	<i>NZD</i>	<i>SEK</i>
<i>Efficient Kitchen-Sink Regression</i>	1.396 ^b	0.947 ^a	0.725	1.365 ^b	0.960 ^b	0.626 ^a	1.286 ^b	0.787 ^a	1.267 ^c
<i>Ridge Regression</i>	-0.032	0.019 ^a	0.180	0.033	0.285 ^a	0.073 ^b	-0.023	0.050	0.006
<i>Lasso Regression</i>	0.118 ^a	0.046	-0.037	0.217 ^b	-0.023	0.728	0.049	0.000	-0.062
<i>Plain Kitchen-Sink Regression</i>	-5.697	-7.265	-8.884	-15.004	-7.028	-7.708	-6.257	-6.074	-8.160
<i>Uncovered Interest Parity</i>	-0.773	0.126 ^a	-2.695	-3.889	-3.723	-0.576	-2.597	-1.129	-5.830
<i>Purchasing Power Parity</i>	-1.664	-1.572	-1.140	-0.961	0.596 ^b	-0.420	-0.756	-1.750	0.262 ^a
<i>Monetary Fundamentals</i>	-0.157	-1.068	-1.556	-1.729	0.242	-1.757	-1.204	-0.912	0.433 ^a
<i>Taylor Rule</i>	-1.763	-1.606	-2.345	-2.635	-2.882	-1.193	-2.078	-1.434	-1.604
<i>Mean Combination</i>	0.480	0.259	0.442	-0.073	0.819 ^b	0.638	-0.069	0.000	-0.305
<i>Median Combination</i>	0.118	0.208	0.361 ^a	0.325 ^a	0.036	0.611 ^a	-0.241	0.585 ^b	0.261
<i>Trimmed Mean Combination</i>	0.222	0.108	0.346	0.254	0.611 ^b	0.737 ^a	-0.094	0.321	0.237
<i>DMSE Combination</i> ($\varphi = 0.90$)	0.652	0.401	0.183	-0.157	0.990 ^b	0.684 ^a	-0.123	-0.046	-0.330
<i>DMSE Combination</i> ($\varphi = 0.95$)	0.578	0.308	0.252	-0.110	1.015 ^b	0.614	-0.104	-0.034	-0.272
<i>DMSE Combination</i> ($\varphi = 1.00$)	0.472	0.258	0.326	-0.089	0.985 ^b	0.530	-0.064	0.013	-0.275
<i>MSE Combination</i> ($\kappa = 60$)	0.490	0.252	0.313	-0.083	0.959 ^b	0.505	-0.047	0.032	-0.266
<i>MSE Combination</i> ($\kappa = 36$)	0.491	0.250	0.288	-0.093	0.923 ^b	0.472	-0.041	0.035	-0.271
<i>MSE Combination</i> ($\kappa = 12$)	0.542	0.170	0.377	0.015	0.801 ^a	0.442	-0.028	0.081	-0.193

Table 4. Testing for Forecast Breakdowns

The table reports the Giacomini and Rossi (2009) t -statistic for testing for forecast breakdowns. This is a test for the stability of the forecasting ability of a model, where the null is that the out-of-sample MSE of the model is equal to the in-sample MSE. The out-of-sample monthly forecasts are obtained in two ways: with recursive regressions that generate forecasts for the period of January 1987 to June 2012 by successively re-estimating the model parameters every time a new observation is added to the sample (Panel A); and with rolling regressions using a rolling window of 10 years that generate forecasts for the same out-of-sample period (Panel B). The superscripts a , b , and c denote statistical significance at the 10%, 5%, and 1% level, respectively, for a one-sided test.

Panel A: <i>Giacomini and Rossi (2009) t-statistics – Recursive Regressions</i>									
	<i>AUD</i>	<i>CAD</i>	<i>CHF</i>	<i>EUR</i>	<i>GBP</i>	<i>JPY</i>	<i>NOK</i>	<i>NZD</i>	<i>SEK</i>
<i>Efficient Kitchen-Sink Regression</i>	1.283 ^a	2.601 ^c	-3.632	-3.193	-3.131	-2.576	0.007	-0.428	0.106
<i>Ridge Regression</i>	1.415 ^a	2.704 ^c	-3.468	-3.016	-2.993	-2.429	0.172	-0.286	0.259
<i>Lasso Regression</i>	1.414 ^a	2.704 ^c	-3.478	-3.021	-2.997	-2.438	0.167	-0.298	0.257
<i>Plain Kitchen-Sink Regression</i>	2.030 ^b	3.292 ^c	-2.282	-1.962	-1.716	-1.276	1.106	0.524	1.012
<i>Uncovered Interest Parity</i>	1.954 ^b	3.185 ^c	-2.609	-2.214	-1.998	-1.441	0.889	0.470	0.886
<i>Purchasing Power Parity</i>	1.966 ^b	3.146 ^c	-2.733	-2.252	-2.270	-1.779	0.874	0.280	0.918
<i>Monetary Fundamentals</i>	1.982 ^b	3.139 ^c	-2.588	-2.210	-2.326	-1.754	0.916	0.299	0.941
<i>Taylor Rule</i>	1.952 ^b	3.143 ^c	-2.792	-2.234	-2.376	-1.823	0.999	0.298	0.879
Panel B: <i>Giacomini and Rossi (2009) t-statistics – Rolling Regressions</i>									
<i>Efficient Kitchen-Sink Regression</i>	0.707	2.874 ^c	-2.450	-3.095	-3.417	-3.393	-0.587	-1.224	0.100
<i>Ridge Regression</i>	0.965	3.073 ^c	-2.130	-2.738	-3.125	-3.105	-0.268	-0.953	0.395
<i>Lasso Regression</i>	0.961	3.078 ^c	-2.121	-2.701	-2.989	-3.116	-0.263	-0.975	0.386
<i>Plain Kitchen-Sink Regression</i>	2.806 ^c	4.468 ^c	0.525	0.133	0.013	-0.876	2.532 ^c	1.111	2.580 ^c
<i>Uncovered Interest Parity</i>	2.183 ^b	4.023 ^c	-0.386	-0.969	-1.346	-1.506	1.225	0.610	1.971 ^b
<i>Purchasing Power Parity</i>	2.109 ^b	3.929 ^c	-0.514	-0.949	-1.556	-1.652	1.319 ^a	0.236	1.815 ^b
<i>Monetary Fundamentals</i>	2.327 ^c	3.954 ^c	-0.373	-0.879	-1.579	-1.769	1.472 ^a	0.417	1.908 ^b
<i>Taylor Rule</i>	2.134 ^b	4.032 ^c	-0.686	-1.227	-1.816	-1.861	1.617 ^a	0.259	1.589 ^a

Table 5. Economic Evaluation of Exchange Rate Predictability

The table shows the out-of-sample economic value of a set of empirical exchange rate models and combined forecasts. We form monthly exchange rate forecasts for nine nominal spot exchange rates relative to the US dollar obtained in two ways: with recursive regressions that generate forecasts for the period of January 1987 to June 2012 by successively re-estimating the model parameters every time a new observation is added to the sample (Panel A); and with rolling regressions using a rolling window of 10 years that generate forecasts for the same out-of-sample period (Panel B). The combined forecasts use combinations of five models: the random walk, uncovered interest parity, purchasing power parity, monetary fundamentals and the Taylor rule. Using these forecasts, we build a maximum expected return strategy subject to a target portfolio volatility $\sigma_p^* = 10\%$ for a US investor who each month dynamically rebalances her portfolio investing in a domestic US bond and nine foreign bonds. For each portfolio, we report the annualized % mean (μ_p), annualized % volatility (σ_p), annualized Sharpe ratio (SR), maximum drawdown (MDD), annualized performance fee (\mathcal{P}) a risk-averse investor is willing to pay to switch from the benchmark random walk strategy to a competing strategy, and the break-even proportional transaction cost (τ^{be}) that cancels out the advantage of a given strategy relative to the random walk. SR_τ and \mathcal{P}_τ account for transaction costs. \mathcal{P} and \mathcal{P}_τ are expressed in annual basis points. τ^{be} is only reported for positive performance measures and is expressed in monthly basis points. The out-of-sample analysis runs from January 1987 to June 2012. The superscripts *a*, *b*, and *c* denote statistical significance at the 10%, 5%, and 1% level, respectively, using the Ledoit and Wolf (2008) bootstrap two-sided test of whether the Sharpe ratio of a model is different from that of the benchmark.

Panel A: <i>Recursive Regressions</i>									
<i>Strategy</i>	μ_p (%)	σ_p (%)	SR	SR_τ	\mathcal{P} (bps)	\mathcal{P}_τ (bps)	MDD (%)	τ^{be} (bps)	
<i>Random Walk</i>	10.7	11.5	0.56	0.38			22.9		
<i>Efficient Kitchen-Sink Regression</i>	15.7	12.2	0.94 ^b	0.79 ^b	414	465	9.4	1526	
<i>Ridge Regression</i>	10.7	11.5	0.56	0.38	6	6	22.6	54	
<i>Lasso Regression</i>	10.6	11.5	0.55	0.37	-6	-6	24.9	-	
<i>Plain Kitchen-Sink Regression</i>	8.4	11.2	0.37	0.16	-215	-249	33.3	-	
<i>Uncovered Interest Parity</i>	11.2	11.0	0.63	0.45	90	89	13.4	44	
<i>Purchasing Power Parity</i>	12.7	12.4	0.68	0.50	126	107	43.6	41	
<i>Monetary Fundamentals</i>	10.4	11.6	0.53	0.34	-36	-53	52.0	-	
<i>Taylor Rule</i>	7.9	11.9	0.30	0.07 ^b	-331	-384	29.6	-	
<i>Momentum</i>	6.6	8.6	0.27	0.04	-234	-226	18.7	-	
<i>1/N strategy</i>	6.1	7.2	0.24	0.24	-230	-36	9.4	-	
<i>Mean Combination</i>	11.8	11.5	0.65	0.46	95	83	34.2	15	
<i>Median Combination</i>	12.0	11.1	0.70	0.50	161	145	16.6	19	
<i>Trimmed Mean Combination</i>	11.7	11.4	0.65	0.45	107	83	24.1	15	
<i>DMSE Combination ($\varphi = 0.90$)</i>	12.0	11.5	0.67	0.47	117	103	32.9	16	
<i>DMSE Combination ($\varphi = 0.95$)</i>	11.9	11.5	0.66	0.46	105	92	33.8	16	
<i>DMSE Combination ($\varphi = 1.00$)</i>	11.8	11.5	0.65	0.46	97	84	34.1	16	
<i>MSE Combination ($\kappa = 60$)</i>	11.8	11.5	0.65	0.46	93	80	34.4	15	
<i>MSE Combination ($\kappa = 36$)</i>	11.7	11.5	0.65	0.45	89	76	34.6	15	
<i>MSE Combination ($\kappa = 12$)</i>	11.6	11.5	0.64	0.44	78	66	34.2	13	

Panel B: <i>Rolling Regressions</i>								
<i>Strategy</i>	μ_p (%)	σ_p (%)	\mathcal{SR}	\mathcal{SR}_τ	\mathcal{P} (bps)	\mathcal{P}_τ (bps)	MDD (%)	τ^{be} (bps)
<i>Random Walk</i>	11.5	13.2	0.55	0.40			50.0	
<i>Efficient Kitchen-Sink Regression</i>	15.8	12.2	0.95 ^b	0.79 ^a	546	521	8.9	522
<i>Ridge Regression</i>	11.1	13.2	0.52	0.37	-35	-40	50.1	-
<i>Lasso Regression</i>	10.2	13.3	0.45 ^b	0.28 ^b	-133	-155	50.0	-
<i>Plain Kitchen-Sink Regression</i>	7.1	11.3	0.25	0.02	-224	-356	47.4	-
<i>Uncovered Interest Parity</i>	9.7	12.7	0.43	0.27	-134	-139	15.3	-
<i>Purchasing Power Parity</i>	8.2	12.7	0.31	0.15	-282	-295	61.4	-
<i>Monetary Fundamentals</i>	9.4	11.8	0.43	0.24	-102	-132	49.4	-
<i>Taylor Rule</i>	8.9	12.4	0.37	0.14	-196	-295	18.3	-
<i>Momentum</i>	6.6	8.6	0.27	0.04	-177	-178	18.7	-
<i>1/N strategy</i>	6.1	7.2	0.24	0.24	-177	12	9.4	-
<i>Mean Combination</i>	11.9	13.1	0.58	0.42	72	29	48.4	4
<i>Median Combination</i>	11.7	12.6	0.59	0.42	109	48	15.2	2
<i>Trimmed Mean Combination</i>	12.3	13.1	0.61	0.45	106	68	54.1	8
<i>DMSE Combination ($\varphi = 0.90$)</i>	11.9	13.0	0.58	0.41	71	26	44.1	3
<i>DMSE Combination ($\varphi = 0.95$)</i>	11.8	13.0	0.58	0.41	70	26	46.5	3
<i>DMSE Combination ($\varphi = 1.00$)</i>	11.8	13.0	0.58	0.41	68	25	48.5	3
<i>MSE Combination ($\kappa = 60$)</i>	11.8	13.0	0.58	0.41	65	22	48.2	3
<i>MSE Combination ($\kappa = 36$)</i>	11.8	13.1	0.58	0.41	60	17	48.3	3
<i>MSE Combination ($\kappa = 12$)</i>	11.8	13.2	0.57	0.40	45	4	48.1	3

Table 6. Long-Horizon Predictive Regressions

The table displays the out-of-sample R_{oos}^2 test statistic in percent for a set of empirical exchange rate models and combined forecasts against the null of a random walk (RW) for predictive horizons of 3-months, 6-months and 12-months ahead. The out-of-sample monthly forecasts are obtained with recursive regressions that generate forecasts for the period of January 1987 to June 2012 by successively re-estimating the model parameters every time a new observation is added to the sample. The combined forecasts use combinations of five models: the random walk, uncovered interest parity, purchasing power parity, monetary fundamentals and the Taylor rule. a , b , and c denote statistical significance at the 10%, 5%, and 1% level, respectively. The superscripts a , b , and c use the Clark and West (2006, 2007) one-sided t -statistic, whereas the subscripts a , b , and c use the Giacomini and White (2006) two-sided t -statistic.

	Panel A: 3-month ahead prediction								
	R_{oos}^2 (%) – Recursive Regressions								
	<i>AUD</i>	<i>CAD</i>	<i>CHF</i>	<i>EUR</i>	<i>GBP</i>	<i>JPY</i>	<i>NOK</i>	<i>NZD</i>	<i>SEK</i>
<i>Efficient Kitchen-Sink Regression</i>	2.455 ^c	2.288 ^c	0.968 ^a	1.461 ^b	−0.526	0.230 ^b	1.633 ^c	2.904 ^c	1.823 ^c
<i>Ridge Regression</i>	0.040 ^b	0.092 ^a	−0.338	−0.206	−0.841	1.031 ^b	−0.020	0.028	−0.054
<i>Lasso Regression</i>	0.057 ^a	−0.225	−1.332	−1.450	−1.851	1.183 ^b	−0.008	−0.002	−0.125
<i>Plain Kitchen-Sink Regression</i>	−8.478	−4.593	−11.302	−12.340	−8.497	−5.966	−5.107	−6.360	−6.954
<i>Uncovered Interest Parity</i>	0.022	−0.315	−7.229	−5.954	−9.591	−0.361	−1.918	−0.366	−3.472
<i>Purchasing Power Parity</i>	−0.251	−1.332	3.088 ^a	2.590 ^b	5.490 ^c	2.525 ^b	1.773 ^a	−1.513	2.788 ^c
<i>Monetary Fundamentals</i>	−0.805	−4.284	−1.949	−1.441	−0.587	−4.588	−2.469	−0.519	−0.605
<i>Taylor Rule</i>	−3.469	−0.863	−3.705	−6.267	−5.313	−2.224	−1.163	−2.463	−1.732
<i>Mean Combination</i>	−0.386	−0.431	1.598 ^c	−0.113	0.475 ^a	2.588 ^c	−0.239	−0.419	−0.072
<i>Median Combination</i>	0.242	0.263	0.775 ^b	1.063 ^c	0.060	3.391 ^c	0.263 ^a	−0.655	0.028
<i>Trimmed Mean Combination</i>	0.230	0.264	0.917 ^b	0.321	0.421 ^a	2.896 ^c	0.181	−0.330	−0.041
<i>DMSE Combination</i> ($\varphi = 0.90$)	−0.142	0.007	2.294 ^c	0.234	0.981 ^b	3.379 ^c	−0.091	−0.138	−0.009
<i>DMSE Combination</i> ($\varphi = 0.95$)	−0.228	−0.160	2.012 ^b	0.076	0.910 ^b	2.945 ^c	−0.140	−0.183	0.002
<i>DMSE Combination</i> ($\varphi = 1.00$)	−0.361	−0.379	1.679 ^c	−0.075	0.646 ^a	2.504 ^c	−0.203	−0.306	−0.041
<i>MSE Combination</i> ($\kappa = 60$)	−0.343	−0.340	1.789 ^b	−0.020	0.676 ^a	2.473 ^c	−0.141	−0.271	−0.030
<i>MSE Combination</i> ($\kappa = 36$)	−0.311	−0.345	1.651 ^b	−0.080	0.781 ^b	2.447 ^b	−0.130	−0.234	−0.041
<i>MSE Combination</i> ($\kappa = 12$)	−0.093	−0.186	2.829 ^c	0.503	0.783 ^b	3.072 ^c	−0.077	−0.057	0.118

Panel B: 6-month ahead prediction									
R_{oos}^2 (%) – Recursive Regressions									
	<i>AUD</i>	<i>CAD</i>	<i>CHF</i>	<i>EUR</i>	<i>GBP</i>	<i>JPY</i>	<i>NOK</i>	<i>NZD</i>	<i>SEK</i>
<i>Efficient Kitchen-Sink Regression</i>	4.898 ^c	3.842 ^c	6.289 ^b	-0.850 ^a	-5.428	4.024 ^c	2.524 ^b	6.438 ^c	1.493 ^c
<i>Ridge Regression</i>	0.049 ^c	-0.602	-3.486	-1.936	-6.516	4.058 ^c	-0.412	0.235	-1.181
<i>Lasso Regression</i>	0.218	-1.274	-4.889	-6.314	-7.759	4.236 ^c	-1.587	0.190	-2.425
<i>Plain Kitchen-Sink Regression</i>	-17.878	-10.225	-17.972	-15.531	-19.635	0.036 ^c	-19.359	-16.664	-11.407
<i>Uncovered Interest Parity</i>	0.322	-0.595	-13.916	-11.926	-15.271	2.233 ^c	-4.506	0.311	-8.215
<i>Purchasing Power Parity</i>	-0.787	-1.389	6.269 ^b	6.080 ^c	12.716 ^c	5.032 ^c	3.859 ^b	-2.553	6.530 ^c
<i>Monetary Fundamentals</i>	-1.851	-9.083	-5.857	-3.113	-1.553	-11.706	-5.239	-0.900	-0.588
<i>Taylor Rule</i>	-7.199	-1.786	-6.535	-7.993	-7.349	-1.668	-3.532	-4.004	-3.725
<i>Mean Combination</i>	-1.032	-1.247	4.327 ^c	2.342 ^b	2.340 ^b	7.571 ^c	-0.611	-0.307	0.311
<i>Median Combination</i>	-0.263	-0.185	2.926 ^c	2.883 ^c	0.722 ^b	8.386 ^c	0.137	-0.461	0.033
<i>Trimmed Mean Combination</i>	0.078	0.329	2.697 ^c	2.509 ^c	2.504 ^c	7.358 ^c	-0.062	-0.712	0.373
<i>DMSE Combination</i> ($\varphi = 0.90$)	-0.304	-0.121	7.085 ^c	3.701 ^c	4.727 ^c	9.603 ^c	0.382	1.546	1.073 ^b
<i>DMSE Combination</i> ($\varphi = 0.95$)	-0.512	-0.465	6.144 ^c	3.238 ^c	4.419 ^c	8.794 ^c	0.214	1.401	1.004 ^b
<i>DMSE Combination</i> ($\varphi = 1.00$)	-0.961	-1.103	4.990 ^c	2.547 ^c	3.433 ^c	7.738 ^c	-0.059	1.035	0.598 ^a
<i>MSE Combination</i> ($\kappa = 60$)	-0.898	-0.939	5.552 ^c	2.831 ^c	3.623 ^c	7.473 ^c	0.205	1.155	0.669 ^a
<i>MSE Combination</i> ($\kappa = 36$)	-0.784	-0.928	5.191 ^c	2.937 ^c	3.846 ^c	7.680 ^c	0.306	1.225	0.751 ^b
<i>MSE Combination</i> ($\kappa = 12$)	-0.184	-0.235	9.648 ^c	5.087 ^c	4.633 ^c	10.965 ^c	0.572	1.865	1.456 ^c

Panel C: 12-month ahead prediction									
R_{OOS}^2 (%) – Recursive Regressions									
	AUD	CAD	CHF	EUR	GBP	JPY	NOK	NZD	SEK
<i>Efficient Kitchen-Sink Regression</i>	-0.029	-6.274	3.901 ^c	8.956 ^c	8.026 ^c	27.850 ^c	-2.249	-2.095	-2.050
<i>Ridge Regression</i>	-3.088	-8.854	2.249 ^c	9.735 ^c	7.221 ^c	27.643 ^c	-3.601	-1.600	-2.799
<i>Lasso Regression</i>	-3.170	-6.109	-3.175	2.611 ^c	16.239 ^c	18.069 ^c	-2.877	-4.538	-7.144
<i>Plain Kitchen-Sink Regression</i>	-32.330	-22.689	-35.049	-12.851 ^c	8.372 ^c	0.257 ^c	-16.359	-13.972	-20.389
<i>Uncovered Interest Parity</i>	2.320 ^c	-1.862	-28.223	-20.073	-19.491	10.042 ^b	-9.661	-5.451	-20.083
<i>Purchasing Power Parity</i>	-1.599	-4.355	11.546 ^c	14.371 ^c	26.898 ^c	13.126 ^c	8.151 ^c	-0.886 ^b	13.997 ^c
<i>Monetary Fundamentals</i>	-3.008	-17.507	-21.501	-13.182	-12.752	-29.599	-9.354	-2.211	-1.535
<i>Taylor Rule</i>	-6.704	-4.661	-7.262	-6.171 ^b	-4.267	-1.248 ^c	-2.317	-4.655	-3.553
<i>Mean Combination</i>	0.455	-2.528	8.600 ^c	8.041 ^c	6.505 ^c	19.329 ^c	-0.004	-0.422	1.918 ^b
<i>Median Combination</i>	-0.300	-0.325	3.252 ^c	5.889 ^c	1.297 ^a	19.279 ^c	0.839 ^a	-0.146	3.018 ^c
<i>Trimmed Mean Combination</i>	1.093 ^b	-0.083	5.612 ^c	6.766 ^c	5.286 ^c	17.682 ^c	1.791 ^b	-0.192	2.666 ^c
<i>DMSE Combination</i> ($\varphi = 0.90$)	2.319 ^b	-0.282	13.547 ^c	11.323 ^c	12.869 ^c	26.579 ^c	1.739 ^b	0.843 ^b	3.920 ^c
<i>DMSE Combination</i> ($\varphi = 0.95$)	1.921 ^b	-0.716	13.059 ^c	10.499 ^c	11.994 ^c	25.327 ^c	1.457 ^b	0.604 ^a	3.708 ^c
<i>DMSE Combination</i> ($\varphi = 1.00$)	0.650 ^a	-1.938	12.911 ^c	9.655 ^c	10.092 ^c	21.693 ^c	0.918 ^a	-0.052	3.330 ^c
<i>MSE Combination</i> ($\kappa = 60$)	0.838 ^a	-1.363	14.921 ^c	10.633 ^c	10.213 ^c	21.335 ^c	1.467 ^b	-0.060	3.825 ^c
<i>MSE Combination</i> ($\kappa = 36$)	1.247 ^b	-0.972	14.096 ^c	11.367 ^c	10.307 ^c	23.961 ^c	1.974 ^c	0.168	4.418 ^c
<i>MSE Combination</i> ($\kappa = 12$)	2.980 ^c	1.093 ^b	25.494 ^c	18.312 ^c	13.561 ^c	37.114 ^c	3.588 ^c	1.908 ^a	6.868 ^c

Table 7. Quarterly Returns

The table displays the out-of-sample R_{oos}^2 test statistic in percent for a set of empirical exchange rate models and combined forecasts against the null of a random walk (RW) using quarterly returns. The out-of-sample quarterly forecasts are obtained with recursive regressions that generate forecasts for the period of January 1987 to June 2012 by successively re-estimating the model parameters every time a new observation is added to the sample. The combined forecasts use combinations of five models: the random walk, uncovered interest parity, purchasing power parity, monetary fundamentals and the Taylor rule. a , b , and c denote statistical significance at the 10%, 5%, and 1% level, respectively. The superscripts a , b , and c use the Clark and (2006, 2007) one-sided t -statistic, whereas the subscripts a , b , and c use the Giacomini and White (2006) two-sided t -statistic.

	R_{oos}^2 (%) – Recursive Regressions								
	<i>AUD</i>	<i>CAD</i>	<i>CHF</i>	<i>EUR</i>	<i>GBP</i>	<i>JPY</i>	<i>NOK</i>	<i>NZD</i>	<i>SEK</i>
<i>Efficient Kitchen-Sink Regression</i>	2.480 _b	2.011 _b	1.672	1.490	1.810 _c	0.591	1.602 _a	3.020 _b	1.875 ^a
<i>Ridge Regression</i>	0.040	0.007	−0.016	−0.023	−1.417	0.248	−0.024	0.042	−0.046
<i>Lasso Regression</i>	0.000	−0.003	−0.165	0.000	0.383	−1.183	0.000	0.000	0.000
<i>Plain Kitchen-Sink Regression</i>	−9.004	−6.620	−13.050	−12.886	−5.151	−6.415	−5.730	−4.068	−7.584
<i>Uncovered Interest Parity</i>	0.211	−0.293	−6.602	−4.292	−7.935	−1.473 ^a	−2.009	0.238	−2.560
<i>Purchasing Power Parity</i>	−0.669	−1.579	3.787 _a	2.911 _a	5.317 _b	2.355 _a	1.782 ^a	−1.741	2.592 _a
<i>Monetary Fundamentals</i>	−0.635	−3.509	−1.247	−1.374	−0.173	−4.103	−1.804	−0.672	−0.568
<i>Taylor Rule</i>	−3.311	−2.743	−2.341	−5.248	−4.988	−1.472	−2.101	−3.614	−2.308
<i>Mean Combination</i>	−0.478	−0.998	2.195 _b	0.269	0.571	2.692 _a	−0.067	−0.392	0.018
<i>Median Combination</i>	0.426	−0.309	−0.213	1.713 _a	0.201	3.611 _b	0.951 ^a	−0.443	0.036
<i>Trimmed Mean Combination</i>	−0.108	−0.040	1.317 ^a	0.866	0.777	2.682 _a	0.673	−0.412	−0.235
<i>DMSE Combination</i> ($\varphi = 0.90$)	−0.487	−0.916	1.435 ^a	−1.356	1.338 ^a	1.858	0.039	−0.406	−0.104
<i>DMSE Combination</i> ($\varphi = 0.95$)	−0.482	−0.946	1.464 ^a	−1.298	1.293 ^a	1.902 ^a	−0.029	−0.428	−0.088
<i>DMSE Combination</i> ($\varphi = 1.00$)	−0.480	−0.988	1.413	−1.282	1.226 ^a	1.856	−0.094	−0.449	−0.099
<i>MSE Combination</i> ($\kappa = 60$)	−0.475	−0.986	1.396	−1.295	1.206 ^a	1.847	−0.088	−0.450	−0.101
<i>MSE Combination</i> ($\kappa = 36$)	−0.479	−0.978	1.481 ^a	−1.300	1.185	1.866	−0.071	−0.467	−0.100
<i>MSE Combination</i> ($\kappa = 12$)	−0.485	−0.988	1.058	−1.536	1.296 ^a	1.452	−0.070	−0.376	−0.164

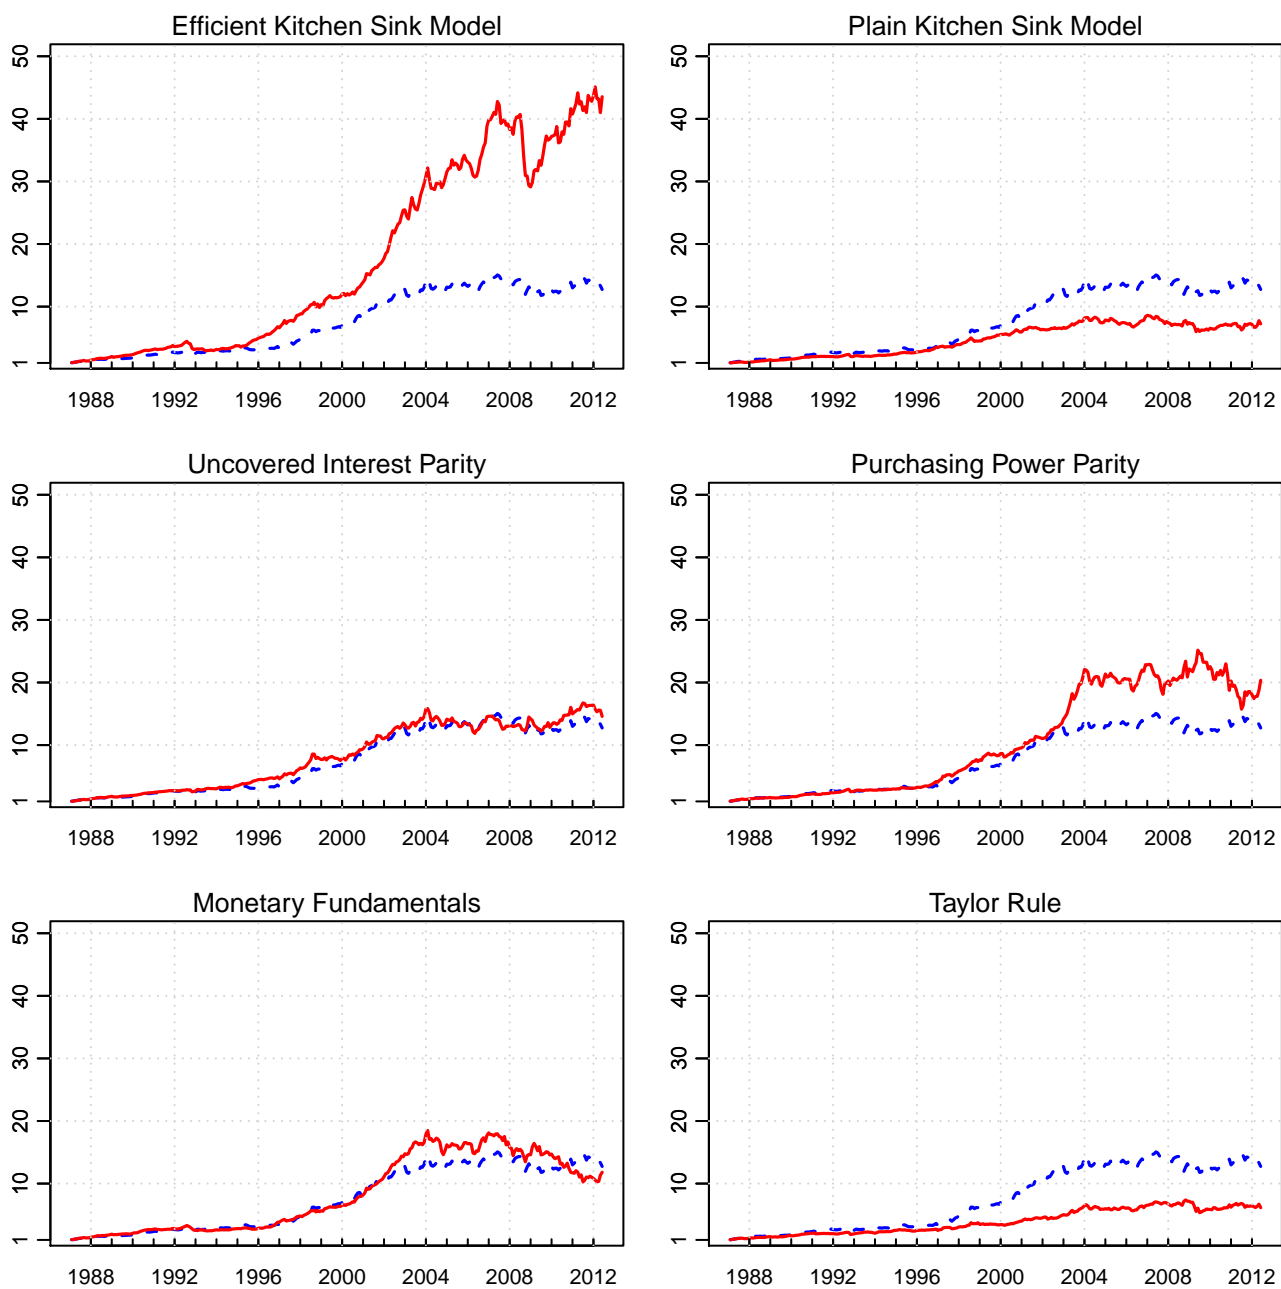


Figure 1. Cumulative Wealth of Empirical Exchange Rate Models

This figure displays the cumulative wealth of out-of-sample dynamic investment strategies conditioning on a set of monthly economic fundamentals (solid red line) relative to the random walk benchmark (dashed blue line). Initial wealth is set at \$1, which grows at the monthly return of the portfolio strategy without transaction costs. The out-of-sample monthly forecasts are obtained with recursive regressions that generate forecasts for the period of January 1987 to June 2012.

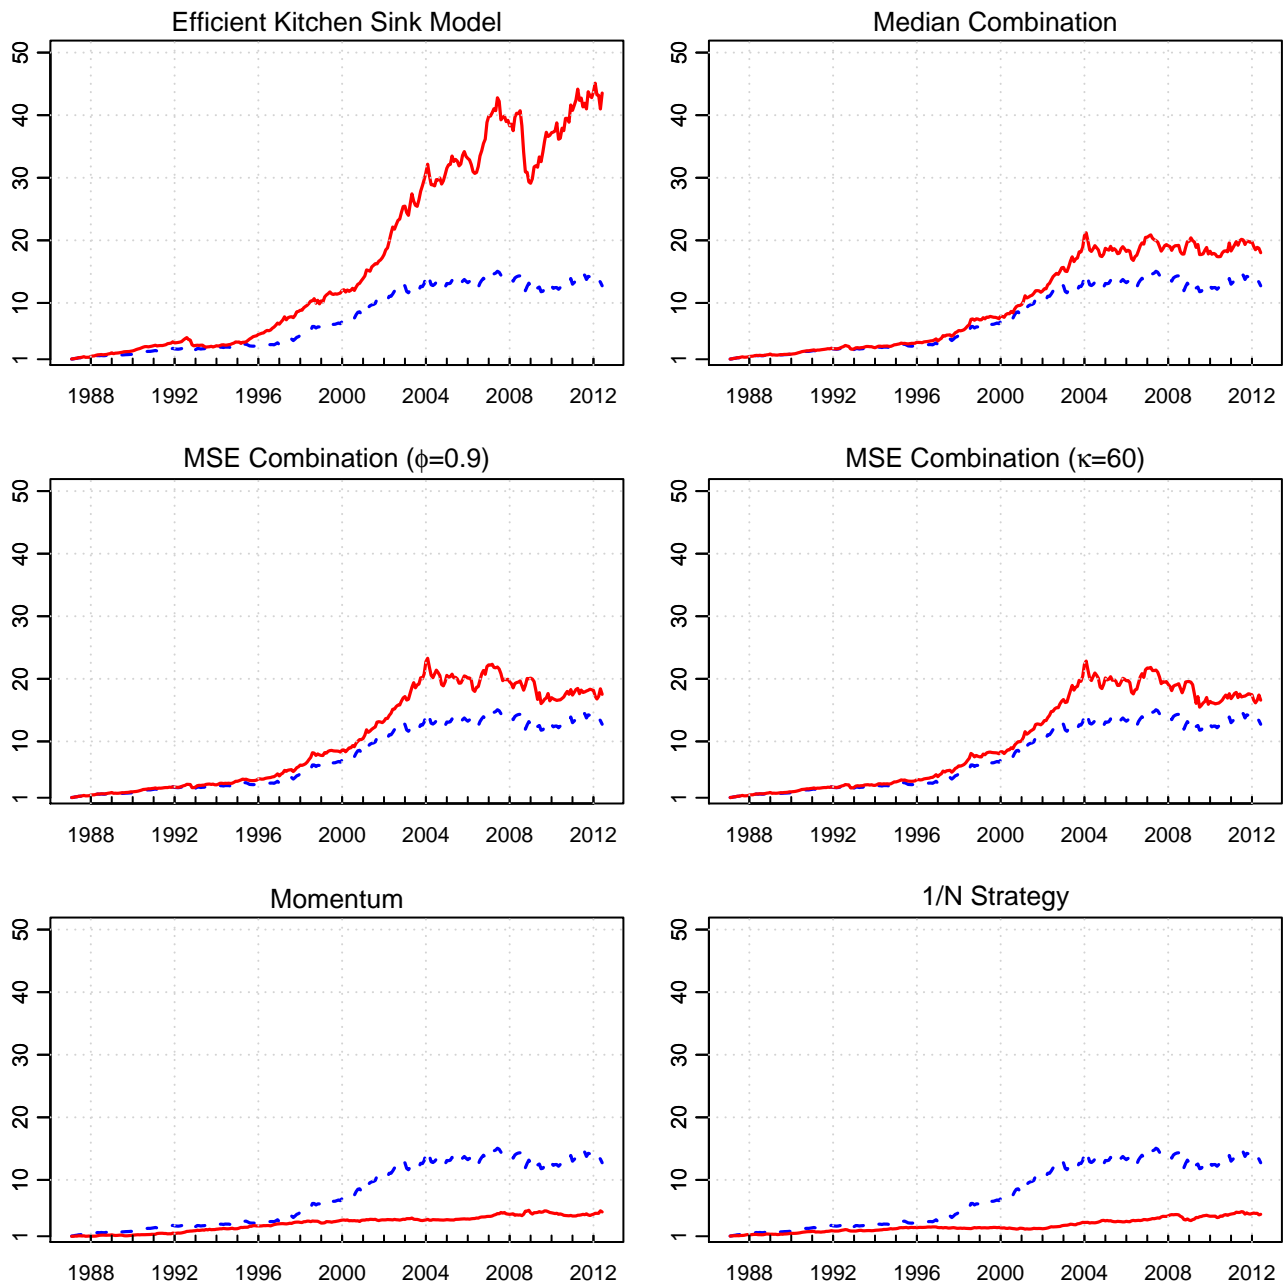


Figure 2. Cumulative Wealth of Combined Forecasting Strategies

This figure displays the cumulative wealth of out-of-sample dynamic investment strategies conditioning on different combined forecasts (solid red line) relative to the random walk benchmark (dashed blue line). Initial wealth is set at \$1, which grows at the monthly return of the portfolio strategy without transaction costs. The out-of-sample monthly forecasts are obtained with recursive regressions that generate forecasts for the period of January 1987 to June 2012.

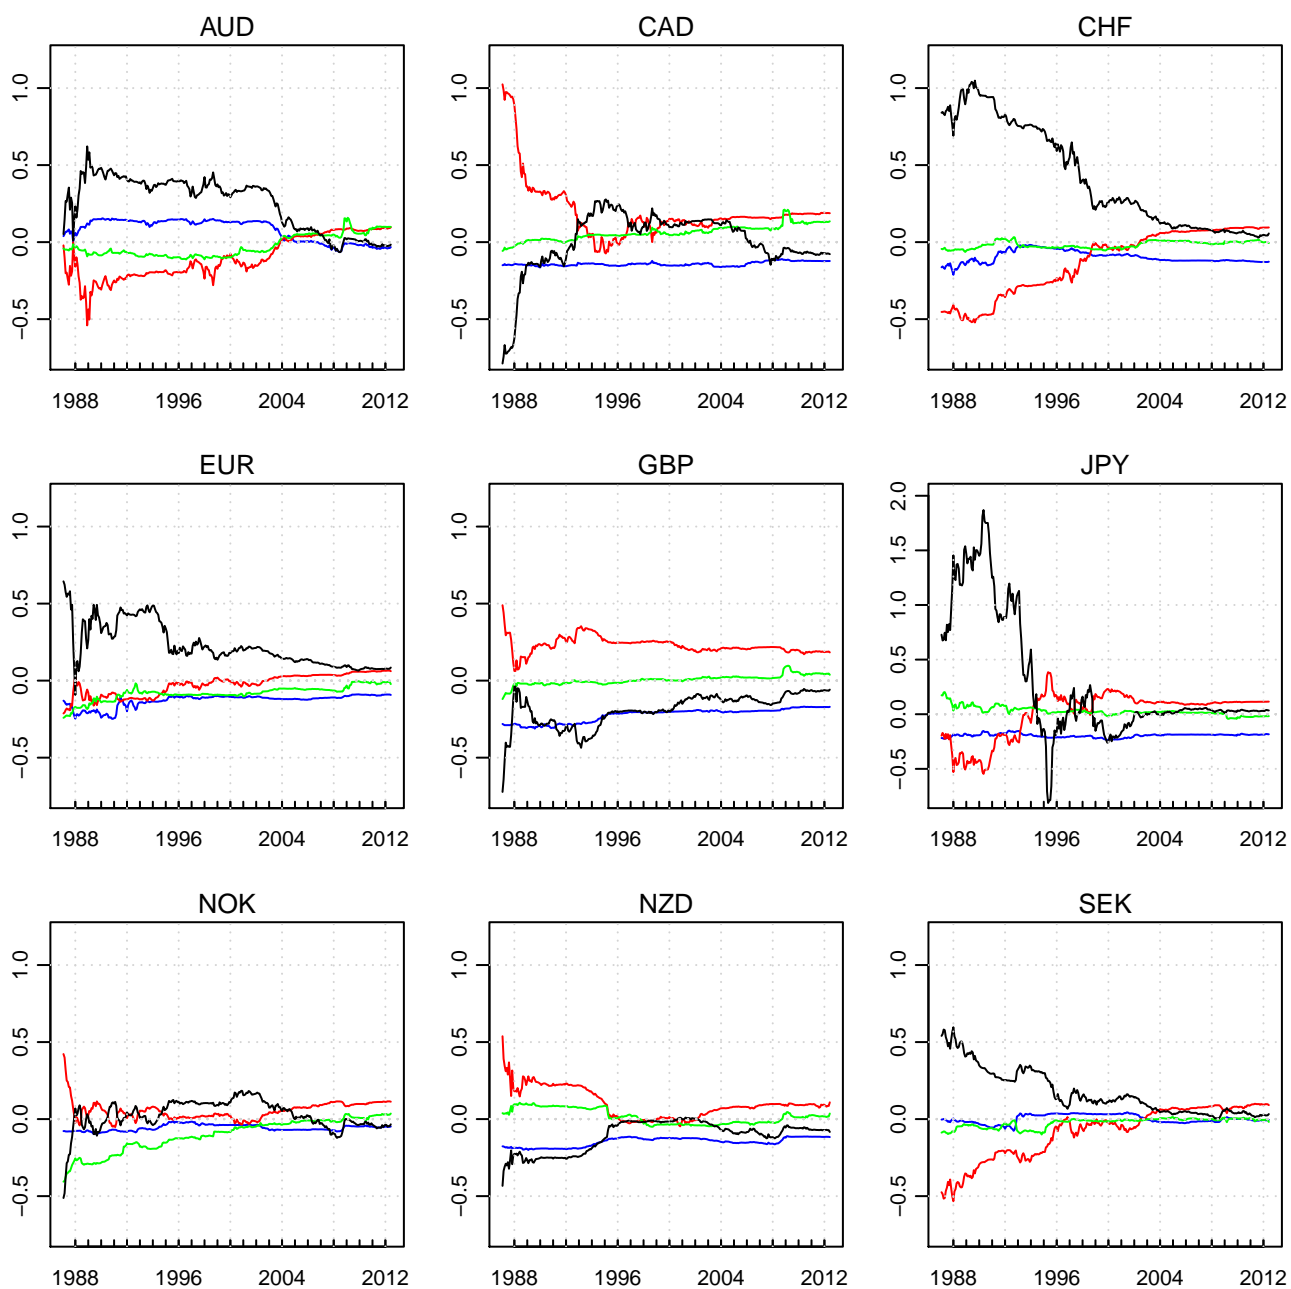


Figure 3. Out-of-Sample Betas for the Plain Kitchen-Sink Regression

This figure displays the out-of-sample variation in the betas of the plain kitchen-sink regression estimated with OLS. The betas are shown for the following regressors: UIP in blue, PPP in red, monetary fundamentals in black and the Taylor rule in green. The out-of-sample monthly forecasts are obtained with recursive regressions that generate forecasts for the period of January 1987 to June 2012. The regressors have been standardized by their sample standard deviation so that each beta can be interpreted as capturing the effect on the exchange rate return of a one standard deviation movement in the regressors. Standardization makes the betas across different regressors directly comparable. All betas are on the same scale except for JPY, where the scale is bigger.

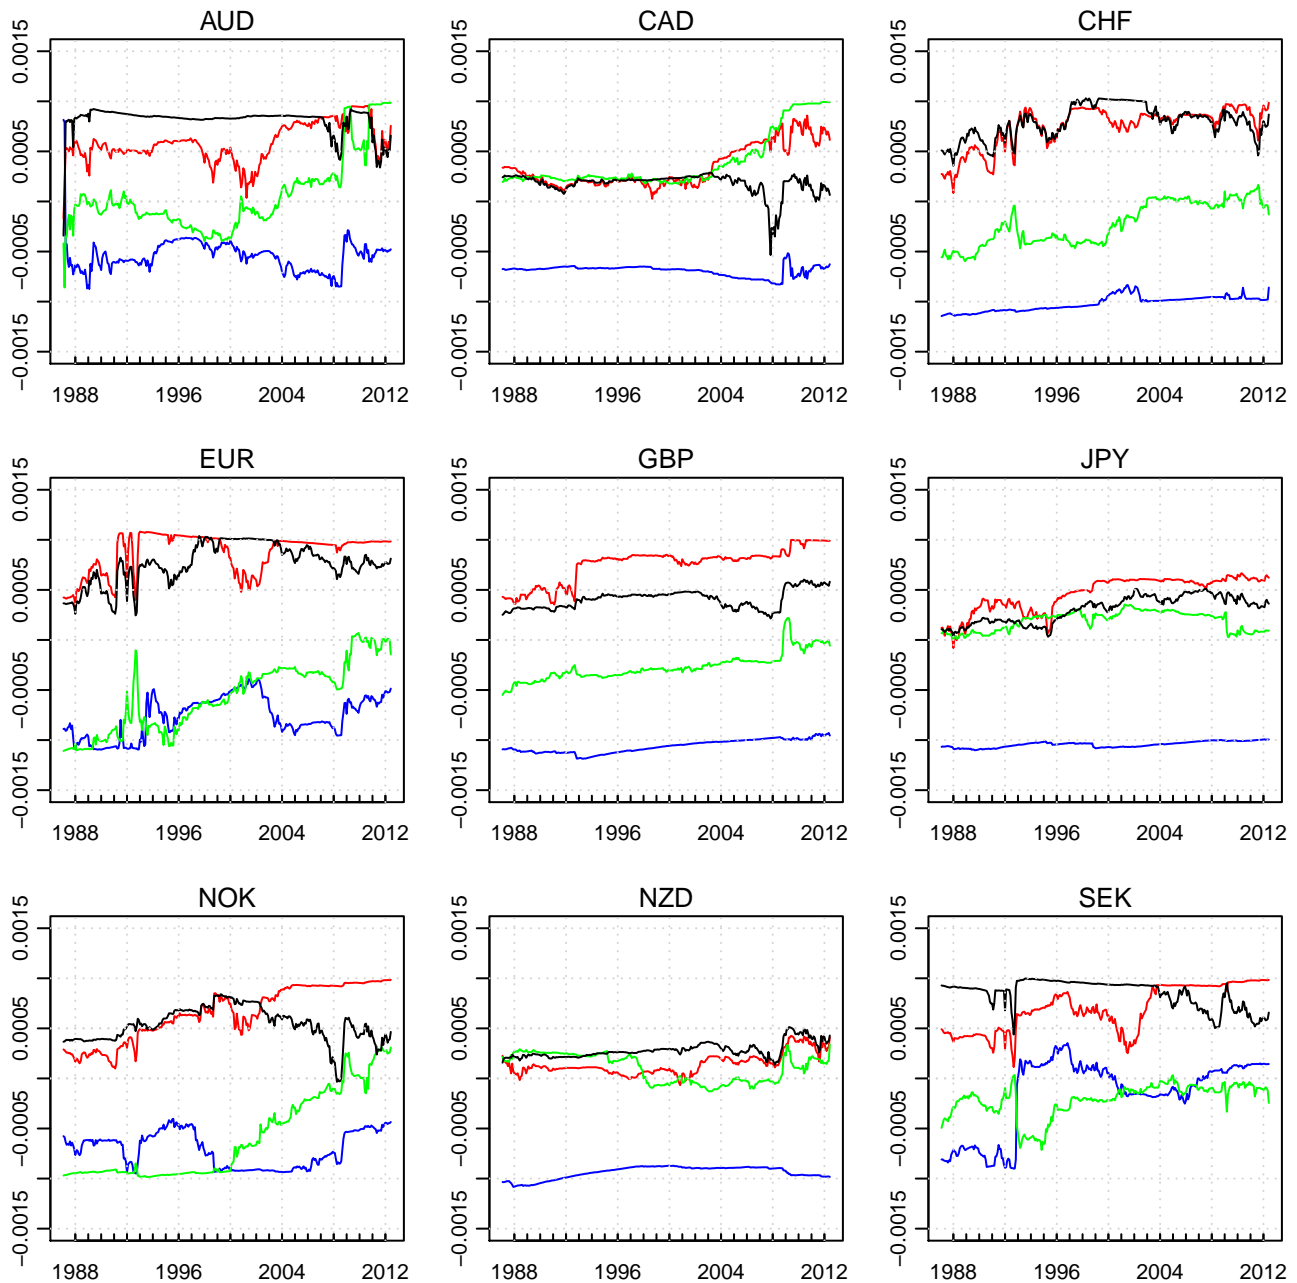


Figure 4. Out-of-Sample Betas for the Efficient Kitchen-Sink Regression

This figure displays the out-of-sample variation in the betas of the efficient kitchen-sink regression estimated with the elastic net. The betas are shown for the following regressors: UIP in blue, PPP in red, monetary fundamentals in black and the Taylor rule in green. The out-of-sample monthly forecasts are obtained with recursive regressions that generate forecasts for the period of January 1987 to June 2012. The regressors have been standardized by their sample standard deviation so that each beta can be interpreted as capturing the effect on the exchange rate return of a one standard deviation movement in the regressors. Standardization makes the betas across different regressors directly comparable. All betas are on the same scale.