


RESEARCH ARTICLE

Exchange rate forecasting and the performance of currency portfolios

Jesus Crespo Cuaresma^{1,2,3,4}  | Ines Fortin⁵ | Jaroslava Hlouskova^{5,6,7}

¹Department of Economics, Vienna University of Economics and Business (WU), Vienna, Austria

²Wittgenstein Centre for Demography and Global Human Capital (WIC), Vienna, Austria

³World Population Program, International Institute for Applied Systems Analysis (IIASA), Laxenburg, Austria

⁴Austrian Institute of Economic Research (WIFO), Vienna, Austria

⁵Research Group Macroeconomics and Economic Policy, Institute for Advanced Studies, Vienna, Austria

⁶Department of Economics, Thompson Rivers University, Kamloops, BC, Canada

⁷Ecosystems Services and Management, International Institute for Applied Systems Analysis, Laxenburg, Austria

Correspondence

Jesus Crespo Cuaresma, Department of Economics, Vienna University of Economics and Business (WU), 1020 Vienna, Austria.
Email: jrcrespo@wu.ac.at

Funding information

Austrian Central Bank, Grant/Award Number: 16250

Abstract

We examine the potential gains of using exchange rate forecast models and forecast combination methods in the management of currency portfolios for three exchange rates: the euro versus the US dollar, the British pound, and the Japanese yen. We use a battery of econometric specifications to evaluate whether optimal currency portfolios implied by trading strategies based on exchange rate forecasts outperform single currencies and the equally weighted portfolio. We assess the differences in profitability of optimal currency portfolios for different types of investor preferences, two trading strategies, mean squared error-based composite forecasts, and different forecast horizons. Our results indicate that there are clear benefits of integrating exchange rate forecasts from state-of-the-art econometric models in currency portfolios. These benefits vary across investor preferences and prediction horizons but are rather similar across trading strategies.

KEYWORDS

currency portfolios, exchange rate forecasting, profitability, trading strategies

1 | INTRODUCTION

Foreign exchange risk is omnipresent in international portfolio diversification, but forecasting exchange rates is well known to be a difficult task. Since the seminal work by Meese and Rogoff (1983), which shows that econometric specifications based on macroeconomic fundamentals are unable to outperform simple random walk forecasts

at short time horizons (up to 1 year), a large number of studies have proposed models aimed at providing accurate out-of-sample predictions of spot exchange rates (see, among others, Berkowitz & Giorgianni, 2001; Boudoukh, Richardson, & Whitelaw, 2008; Cheung, Chinn, & Pascual, 2005; Chinn & Meese, 1995; Kilian, 1999; MacDonald & Taylor, 1994; Mark, 1995; Mark & Sul, 2001). In parallel, a literature has emerged which examines empirically the

potential profitability of technical trading rules based on exchange rate predictions (see Menkhoff & Taylor, 2007, for a review). Although the random walk specification has naturally emerged as the benchmark to beat in terms of out-of-sample predictive accuracy, it is not clear that it will also yield the most profitable trading strategy. Portfolio managers are expected to be more concerned with profitability than with out-of-sample accuracy.

Our study aims at addressing how the joint modeling of exchange rates and fundamentals provide economic value in terms of improving currency portfolio performance. We therefore contribute to the long-standing literature on the use of exchange rate models based on fundamentals for forecasting, taking a new perspective in the evaluation of different econometric specifications. We provide an evaluation framework where we take the perspective of a currency portfolio manager (investor) who follows trading strategies based on exchange rate forecasts and whose main goal is to maximize (risk-adjusted) profits, under certain types of preferences. Our currency portfolio manager considers the exchange rates of the euro against the US dollar (USD), the British pound (GBP), and the Japanese yen (JPY), and for each of these three exchange rates creates a “single asset.” The returns of this asset are implied by a certain trading strategy that is based on exchange rate forecasts. The optimal portfolio is then made up of these three single assets according to the manager’s—or some investor’s—preferences.

The two primary research questions in our study are the following. First, does the information on exchange rate fundamentals provide valuable information to construct optimal currency portfolio that outperform simple benchmark portfolios, and thus is there a value added in engaging in active portfolio management—or can the portfolio manager achieve the same (risk-adjusted) profit by just investing in some simpler assets (benchmark portfolios)? As simpler assets we consider the single assets of which the optimal portfolio consists as well as the equally weighted portfolio based on forecasts from the model based on macroeconomic fundamentals as well as on random walk predictions. This research question links our work to the large literature on the statistical and economic evaluation of exchange rate forecasts (see Abbate & Marcellino, 2018; Della Corte, Sarno, & Tsiakas, 2009; Rossi, 2013) and provides a novel evaluation context that goes beyond the existing methods based on forecast errors and directional change statistics.

Relating to the first question, there is some empirical evidence indicating that simple portfolios, like equally weighted portfolios, are not necessarily outperformed (e.g., in terms of the Sharpe ratio) by more complex portfolios (see DeMiguel, Garlappi, & Uppal, 2009; Jacobs, Müller, & Weber, 2014). The existing evidence in the liter-

ature, however, relates to equity markets (DeMiguel et al. 2009) and equity, bond and commodity markets (Jacobs et al. 2014), and it is not obvious that these findings carry over to foreign exchange markets. Our study contributes to enlarge this body of empirical evidence by concentrating on foreign exchange markets. In order to compare the different currency portfolios, we employ a number of (risk-adjusted) performance measures, including the Omega measure, the Sharpe ratio, and the Sortino ratio. We consider all the multivariate time series models and the methods of forecast combinations entertained in Costantini, Crespo Cuaresma, and Houskova (2014, 2016) to generate exchange rate forecasts.¹

We consider two different trading strategies in constructing the single assets. The first one is the simple “buy low, sell high” trading strategy described, for example, in Gençay (1998), where the trading signal is based on the spot exchange rate and its forecast. The second one is based on exploiting the forward rate unbiased expectation hypothesis, using forward contracts, and is similar to the carry trade strategy used, for example, in Burnside, Eichenbaum, and Rebelo (2008). In this case the trading signal is based on the forward exchange rate and the exchange rate forecast. In order to assess the performance of optimal currency portfolios versus benchmark portfolios (single assets, equally weighted portfolio), we use a data-snooping bias-free test, which is based on an extensive bootstrap procedure. By employing this test we ensure that the performance superiority of certain optimal portfolios—if any—is systematic and not merely due to luck. The test identifies which optimal portfolios significantly beat the benchmark portfolio in terms of certain risk-adjusted performance measures. In addition, we assess the values of optimal portfolios with respect to benchmark portfolios by computing break-even transaction costs (see Della Corte et al., 2009; Della Corte & Tsiakas, 2013).

Returns implied by trading strategies have also been investigated in other exchange rate studies. Burnside et al. (2008), for example, examine the returns implied by the carry trade strategy, which determines to sell (buy) a currency forward when it trades at a forward premium (discount). This trading strategy is similar to our second trading strategy. The authors apply the carry trade strategy to individual currencies as well as to an equally weighted portfolio of 23 currencies and find that constructing a portfolio improves the performance of the carry trade strategy substantially: the Sharpe ratio of the equally weighted carry trade strategy is more than 50% higher than the

¹See also Crespo Cuaresma and Hlouskova (2005), Crespo Cuaresma (2007), Costantini and Pappalardo (2010), and Costantini and Kunst (2011).

median Sharpe ratio across currency-specific carry trade strategies. Unlike our study, Burnside et al. (2008) do not test the statistical significance of the portfolio out-performance with respect to the single-currency-based asset. While they use a simple, equally weighted portfolio, we take the investor's preferences into account explicitly and optimize the portfolio according to these preferences. Another fundamental difference with respect to their work is that we include exchange rate forecasts in the definition of our trading strategies with the aim of improving their performance. Burnside et al. (2008), on the other hand, use bid-and-ask spot and forward exchange rates. In a related paper, Burnside et al. (2011) examine carry trade and momentum strategies for single currencies and equally weighted currency portfolios, and review possible explanations for the profitability of these strategies.

Recently, the work by Barroso and Santa-Clara (2015) aims at maximizing the expected return of a portfolio in the forward exchange market, given preferences described by the power utility and using the parametric portfolio policies approach of Brandt, Santa-Clara, and Valkanov (2009). Papaioannou, Portes, and Siourounis (2006) assess how the introduction of the euro as a new currency has potentially changed the optimal composition of portfolios of foreign exchange reserves and how the optimal holdings compare to actual reserve portfolios held by central banks.

The economic value of using model-based exchange rate forecasts to construct optimal portfolios has several dimensions. First, there is economic value which can be exclusively attributed to portfolio optimization and can be assessed by comparing the performance of optimal portfolios with the performance of the equally weighted portfolio based on composite forecasts. Second, we can consider the economic value of both portfolio optimization and exchange rate forecasting, which can be assessed by comparing the performance of optimal portfolios with the performance of the equally weighted portfolio based on the random walk. Third, we can assess the exclusive economic value of exchange rate forecasting based on fundamentals by comparing the performance of benchmark portfolios based on composite forecasts with that of benchmark portfolios based on the random walk.² Our main results imply that a positive economic value can be observed for forecast horizons of 1, 3, and 6 months, independently of which one of the three dimensions of economic value we consider. This applies for the mean-variance and conditional-value-at-risk optimal portfolios with respect to the equally weighted portfolio. For a forecast horizon of 12 months, however, we only observe a positive economic value of portfolio optimization but not of

forecasting, and also not of portfolio optimization and forecasting.

The paper is structured as follows. In Section 2 we present the analytical framework required for exchange rate forecasting and portfolio optimization. First we introduce the individual exchange rate forecast models and forecast combinations methods that we use to generate the exchange rate predictions. Then we describe how we compute the *best* forecasts (composite forecasts) and present the different types of preferences and the corresponding optimization problems. We conclude Section 2 by describing the risk-adjusted performance measures we consider, as well as the data-snooping bias-free test for equal performance and the calculation of break-even transaction costs. Section 3 discusses the empirical results and Section 4 concludes.

2 | ANALYTICAL FRAMEWORK: EXCHANGE RATE FORECASTING AND OPTIMAL CURRENCY PORTFOLIOS

2.1 | Exchange rate specifications

We start by describing the modeling framework used to obtain ensembles of exchange rate predictions that can be used to construct currency portfolios. The class of specifications we entertain in order to obtain forecasts of the exchange rate can be conceptualized in the context of the so-called monetary model of exchange rates originally developed in the work by, for example, Frenkel (1976), Dornbusch (1976), or Hooper and Morton (1982). The monetary model of exchange rate determination has been often used as a theoretical framework to create exchange rate predictions based on macroeconomic fundamentals (see Costantini et al., 2016; Crespo Cuaresma & Hlouskova, 2005). Starting with standard Cagan money demand equations for the domestic and foreign economy, the monetary model of exchange rate formation assumes that purchasing power parity acts as a long-run equilibrium and thus leads to a relationship between the exchange rate on the one hand and money supply, interest rates, and income levels in the two economies on the other hand.

Consider a money demand equation in the domestic economy given by

$$m_t - p_t = \alpha y_t + \beta i_t, \quad (1)$$

where m_t is the nominal money demand (in logs), p_t denotes the price level (in logs), y_t is a measure of income (in logs) and i_t is the interest rate. Assuming a similar specification in the foreign economy, where asterisks denote the corresponding parameters and variables, we can write

²This issue is the main focus of Costantini et al. (2016).

the real money demand equation as

$$m_t^* - p_t^* = \alpha^* y_t^* + \beta^* i_t^*. \quad (2)$$

The long-run equilibrium condition of the monetary model is given by the purchasing power parity (PPP) condition, equating the nominal exchange rate (S_t) to the price differential between the two economies:

$$\log S_t = p_t - p_t^* = m_t - m_t^* - \alpha y_t + \alpha^* y_t^* - \beta i_t + \beta^* i_t^*. \quad (3)$$

This specification calls for the use of money supply, income, and interest rates as potential covariates in models aimed at assessing exchange rate dynamics. If, in addition, the uncovered interest rate parity (UIP) is assumed, together with identical interest rate semi-elasticities in both economies, the resulting specification will not include interest rates as a potential determinant of exchange rate movements.

Given this modeling framework, the relationship between the exchange rate and its determinants tends to be routinely specified empirically in the form of vector autoregressive (VAR) and vector error correction (VEC) models. Defining the vector z_t , which contains the (log) exchange rate, the (log) money supply in the domestic and foreign economy, the (log) production index in the domestic and foreign economy and the respective short- and long-term interest rates, this implies that its dynamics are given by

$$z_t = \Phi_0 + \sum_{l=1}^p \Phi_l z_{t-l} + \epsilon_t, \quad \epsilon_t \sim \text{n.i.d.}(\mathbf{0}, \Sigma_\epsilon), \quad (4)$$

where Φ_l for $l = 1, \dots, p$ are matrices of coefficients and Φ_0 is a vector of constants. Alternatively, if the variables of the model are linked by one or more cointegration relationships which act as a long-run attractor of the data, the specification given by Equation 4 can be written in vector error correction (VEC) form:

$$\Delta z_t = \Theta_0 + \delta \psi' z_{t-1} + \sum_{l=1}^{p-1} \Theta_l \Delta z_{t-l} + \epsilon_t. \quad (5)$$

In this specification the cointegration relationships are given by $\psi' z_t$ and δ quantifies the speed of adjustment to the long-run equilibrium. Alternatively, if the variables in z_t are unit-root nonstationary but no cointegration relationship exists among them, a VAR model in first differences (DVAR) would be the appropriate representation, which amounts to the model given in Equation 5 with $\delta = 0$. Alternatively to modeling each exchange rate individually, a large vector autoregressive structure where fundamentals for all pairs of countries and their respective exchange rates are included could also be chosen as a modeling framework, but the extremely large number

of parameters involved in such a model would make the estimation difficult.

Myriads of studies have used VAR and VEC models based on macroeconomic variables as specifications for exchange rate predictions, with no robust evidence for improved forecasting ability as compared to the random walk in the short term but often better results for predictions in longer horizons, in particular for VEC specifications. There is some agreement in the literature that forecasts generated by models which explicitly address the potential existence of long-run equilibria in the form of cointegration relationships (VEC specifications) tend to systematically outperform forecasts generated by the naive random walk model in terms of the mean square error at horizons of around one year (see for example MacDonald & Taylor, 1994). This result is, however, far from being homogeneous across currencies and time periods. It should be noted that our contribution uses performance measures that do not correspond to those utilized in the exchange rate forecasting literature hitherto, and as such the existing results may not be a perfect reference for comparison.

We entertain several types of vector autoregressive and vector error correction models as specifications for exchange rate forecasting. On the one hand, we differentiate between restricted and unrestricted models depending on whether the foreign and domestic covariates are included as individual variables in the model or as a single covariate measuring the domestic–foreign difference. We refer to models containing the latter as *restricted* models (r-VAR, r-DVAR, r-VEC), whereas the models based on separated domestic and foreign variables are labeled *unrestricted* models (VAR, DVAR, VEC). We also consider subset-VAR models, where statistically insignificant lags of the variables are omitted, and label them s-VAR, s-DVAR, rs-VAR, and rs-DVAR.

In terms of estimation method, we consider multivariate models estimated using standard frequentist methods and Bayesian VARs. Bayesian DVARs are estimated using the standard Minnesota prior (see Doan, Litterman, & Sims, 1984; Litterman, 1986). The lag length of all multivariate model specifications under consideration is selected using the Akaike information criterion (AIC) for potential lag lengths ranging from 1 to 12 lags. For the VEC models, selection of the lag length and the number of cointegration relationships is carried out simultaneously using the AIC. Table 1 lists the 12 individual forecast models used.

2.2 | Forecast combinations

The set of methods used to create forecast combinations from individual multivariate time series models is similar to that in Costantini et al. (2016). Let S_t^j be the spot

TABLE 1 Individual forecast models

VAR(p)	Vector autoregression in levels based on domestic and foreign variables with p lags
DVAR(p)	Vector autoregression in first differences based on domestic and foreign variables with p lags
VEC(c, p)	Vector error correction model based on domestic and foreign variables with c cointegration relationships
r-VAR(p)	Restricted VAR, based on differences between domestic and foreign variables
r-DVAR(p)	Restricted DVAR, based on differences between domestic and foreign variables
r-VEC(c, p)	Restricted VEC, based on differences between domestic and foreign variables with c cointegration relationships
s-VAR(p)	Subset vector autoregression in levels based on domestic and foreign variables with p lags
s-DVAR(p)	Subset vector autoregression in first differences based on domestic and foreign variables with p lags
rs-VAR(p)	Restricted subset VAR, based on differences between domestic and foreign variables
rs-DVAR(p)	Restricted subset DVAR, based on differences between domestic and foreign variables
BDVAR(p)	Bayesian vector autoregression in first differences based on domestic and foreign variables
r-BDVAR(p)	Bayesian vector autoregression in first differences based on differences between domestic and foreign variables

exchange rate of euros (EUR) per foreign currency unit (FCU) of currency j , that is, EUR/FCU, at time t , and let $\hat{S}_{m,t+h|t}^j$ be the exchange rate forecast of euros per foreign currency unit of currency j obtained using model m , $m = 1, \dots, M$, for time $t + h$ conditional on the information available at time t (i.e., h is the forecast horizon). In the following we drop superscript j to keep the exposition simpler. The combinations of forecasts entertained in this study, $\hat{S}_{c,t+h|t}$, take the form of a linear combination of the predictions of individual specifications:

$$\hat{S}_{c,t+h|t} = w_{c,0t}^h + \sum_{m=1}^M w_{c,mt}^h \hat{S}_{m,t+h|t}, \quad (6)$$

where c is the combination method, M is the number of individual forecasts and the weights are given by $\{w_{c,mt}^h\}_{m=0}^M$.

Since several combination methods require statistics based on a hold-out sample where the relative predictive ability of models is assessed, let us introduce here some notation on subsample limits: T_0 is used to denote the first observation of the available sample, the interval (T_1, T_2) is used as a hold-out sample to obtain weights for those methods where such a subsample is required, and T_3 is the last available observation. The sample given by (T_2, T_3) is the proper out-of-sample period used to compare the different methods.

In order to pool the forecasts of individual specifications, we consider a large number of combination methods proposed in the literature. These are the same methods that have been recently used in Costantini et al. (2016) to evaluate exchange rate predictability:³

- *Mean, trimmed mean, median.* For the mean prediction, $w_{\text{mean},0t}^h = 0$ and $w_{\text{mean},mt}^h = \frac{1}{M}$ in Equation 6. The trimmed mean uses $w_{\text{trim},0t}^h = 0$ and $w_{\text{trim},mt}^h = 0$ for the individual models that generate the smallest and largest

forecasts, while $w_{\text{trim},mt}^h = \frac{1}{M-2}$ for the remaining individual models. For the median combination method, $\hat{S}_{c,t+h|t} = \text{median}\{\hat{S}_{m,t+h|t}\}_{m=1}^M$ is used (see Costantini & Pappalardo, 2010).

- *Ordinary least squares (OLS) combination.* The weights of this method coincide with the estimated coefficients obtained by regressing actual exchange rates on a constant and corresponding exchange rate forecast. In our application we use a rolling window over the hold-out sample. Granger and Ramanathan (1984) provide more details on this simple forecast pooling methodology.
- *Combination based on principal components (PC).* This method allows us to overcome multicollinearity of predictions across models by reducing them to a few principal components (factors). The method is identical to the OLS combining method by replacing forecasts with their principal components. In our application, we choose the number of principal components using the variance proportion criterion, which selects the smallest number of principal components such that a certain fraction of variance is explained. We set the proportion to 80%.⁴
- *Combination based on the discount mean square forecast errors (DMSFE).* Following Stock and Watson (2004), the weights in Equation 6 depend inversely on the historical forecasting performance of individual models, $w_{\text{DMSFE},m,t}^h = \text{WMSE}_{mth}^{-1} / \sum_{l=1}^M \text{WMSE}_{lth}^{-1}$, where $\text{WMSE}_{mth} = \sum_{i=T_1+h-1}^t \theta^{T-h-i} (S_i - \hat{S}_{m,i|\bar{i}-h})^2$, for $t = T_2 - h, \dots, T_3 - h$, $m = 1, \dots, M$, $w_{\text{DMSFE},0,t}^h = 0$, and θ is a discount factor. In the empirical application, we use $\theta = 0.95$.
- *Combination based on hit/success rates (HR).* The method uses the proportion of correctly predicted

³Costantini et al. (2016) provide a more detailed discussion of these forecast averaging methods.

⁴More details on the method are provided in Hlouskova and Wagner (2013), where the principal components augmented regressions are used in the context of the empirical analysis of economic growth differentials across countries. Except for Costantini et al. (2016), we are not aware of the existence of any study using this approach in the context of the exchange rate forecasts.

directions of exchange rate changes of model m to the number of all correctly predicted directions of exchange rate changes by the models used, $w_{HR,mt}^h = \sum_{i=T_1+h-1}^t DA_{m,ih} / \sum_{l=1}^M \left(\sum_{j=T_1+h-1}^t DA_{l,ih} \right)$, where $t = T_2 - h, \dots, T_3 - h$ and the index of directional accuracy is given by $DA_{m,th} = I(\text{sgn}(S_t - S_{t-h}) = \text{sgn}(\hat{S}_{m,t|t-h} - S_{t-h}))$, where $I(\cdot)$ is the indicator function.

- *Combination based on the exponential of hit/success rates (EHR)* (Bacchini, Ciammola, Iannaccone, & Marini, 2010). The weights in this method are $w_{EHR,mt}^h = \exp\left(\sum_{i=T_1+h-1}^t (DA_{m,ih} - 1)\right) / \sum_{l=1}^M \exp\left(\sum_{i=T_1+h-1}^t (DA_{l,ih} - 1)\right)$, where $t = T_2 - h, \dots, T_3 - h$.
- *Combination based on the economic evaluation of directional forecasts (EEDF)*. The weights in this method capture the ability of models to predict the direction of change of the exchange rate, while taking into account the magnitude of the realized change and are thus given by $w_{EEDF,mt}^h = \sum_{i=T_1+h-1}^t DV_{m,ih} / \sum_{l=1}^M \left(\sum_{i=T_1+h-1}^t DV_{l,ih} \right)$, where $t = T_2 - h, \dots, T_3 - h$ and $DV_{m,th} = |S_t - S_{t-h}| DA_{m,th}$.
- *Combination based on predictive Bayesian model averaging (BMA)*. The weights used are based on the corresponding posterior model probabilities based on out-of-sample (rather than in-sample) fit. See, for example, Raftery, Madigan, and Hoeting (1997), Carriero, Kapetanios, and Marcellino (2009), Crespo Cuaresma (2007), and Feldkircher (2012).

We create weights based on comparing log-predictive scores of the different models in a hold-out subsample, as this prediction error statistic is routinely used in methodological comparisons involving BMA (see, e.g., Fernandez, Ley, & Steel, 2001). The weights are thus given by

$$w_{BMA,mt}^h = \frac{\sum_{i=T_1+h-1}^t \exp\left\{-\left[\frac{1}{2} \log\left(2\pi\hat{\sigma}_{m,\tilde{i}|i-h}^2\right) + \frac{1}{2}\left(\frac{S_i - \hat{S}_{m,\tilde{i}|i-h}}{\hat{\sigma}_{m,\tilde{i}|i-h}}\right)^2\right]\right\}}{\sum_{l=1}^M \sum_{i=T_1+h-1}^t \exp\left\{-\left[\frac{1}{2} \log\left(2\pi\hat{\sigma}_{l,\tilde{i}|i-h}^2\right) + \frac{1}{2}\left(\frac{S_i - \hat{S}_{l,\tilde{i}|i-h}}{\hat{\sigma}_{l,\tilde{i}|i-h}}\right)^2\right]\right\}}, \quad (7)$$

where $\hat{\sigma}_{m,\tilde{i}|i-h}^2$ is the estimated variance of the corresponding prediction for model m , $\hat{S}_{m,\tilde{i}|i-h}$, and $t = T_2 - h, \dots, T_3 - h$.

- *Combinations based on frequentist model averaging (FMA)* (see Claeskens & Hjort, 2008; 2003). The weights are calculated as $w_{FMA,mt}^h = \exp\left(-\frac{1}{2}IC_{mt}\right) / \sum_{l=1}^M \exp\left(-\frac{1}{2}IC_{lt}\right)$, where IC_{mt} stands

for an information criterion of model m and t is the last time point of the data over which models are estimated.

We use combinations of forecasts based on the AIC, Schwarz criterion (BIC), and Hannan–Quinn criterion (HQ). The weights corresponding to the BIC can be interpreted as an approximation to the posterior model probabilities in BMA.⁵

A list of all forecast combination methods we use can be found in Table 6 in the supporting information Appendix.

2.3 | Predictive accuracy

Exchange rate forecasts are evaluated using the traditional mean square error (MSE).⁶ We obtain the standard square error $SE_{m,t,h} = (\hat{S}_{m,t|t-h} - S_t)^2$ by a rolling-window estimation; that is, we keep the estimation sample size constant (equal to $T_1 - T_0$) as we re-estimate the models, thus moving the window that defines the sample used to estimate the model parameters. The MSE for each model and forecast combination method is thus calculated over the out-of-sample period for a given forecast horizon h and aggregated as⁷ $MSE_{mh} = 1/(T_3 - T_2 + 1) \sum_{j=0}^{T_3-T_2} SE_{m,T_2+j,h}$.

In addition, we compute composite forecasts based on the MSE of predictions from all models and combination methods over a certain period. In particular, for this technique at each time point t we choose the model or forecast combination method (and thus also the forecast for time point $t + h$) with the minimum MSE over a certain time window ending at time point t , $\hat{S}_{t+h|t}^{MSE,l} = \hat{S}_{m_{lth}^{MSE},t+h|t}$, where $m_{lth}^{MSE} = \arg \min_m \sum_{j=l}^t SE_{m,j,h}$. Time point l , such that $T_2 \leq l \leq t$, defines the beginning of the window over which the performance is evaluated. The evaluation window is $[l, t]$, where $l \leq t \leq T_3$. For our empirical results we use $l = t - 12$; that is, the model or forecast combination method with the smallest MSE over the last 12 months is chosen.

⁵See Raftery et al. (1997) and Sala-i-Martin, Doppelhofer, & Miller (2004).

⁶In a previous version of this paper we also use profit measures such as the directional value, statistics that capture the economic value of directional forecasts and returns with respect to the two trading strategies (see Crespo Cuaresma, Fortin, & Hlouskova, 2017). However, we restrict ourselves to the MSE here.

⁷The term *model* and thus also its abbreviation m is used in a broader sense that includes both individual models and forecasting rules and methods (like forecast combinations and composite forecasts).

$$y_{th}^{jm,TS2} = \begin{cases} 1, & \text{if } \hat{S}_{t+h|t}^{jm} < F_{t+h|t}^j \quad (\text{one FCU of currency } j \text{ is sold forward at } t \text{ and bought at } t+h) \\ -1, & \text{if } \hat{S}_{t+h|t}^{jm} > F_{t+h|t}^j \quad (\text{one FCU of currency } j \text{ is bought forward at } t \text{ and sold at } t+h), \end{cases} \quad (9)$$

$$r_{t+h,h}^{jm,TS2} = \begin{cases} \frac{1}{F_{t+h|t}^j} (F_{t+h|t}^j - S_{t+h}^j) = 1 - \frac{S_{t+h}^j}{F_{t+h|t}^j}, & \text{if } y_{th}^{jm,TS2} = 1 \\ \frac{1}{F_{t+h|t}^j} (S_{t+h}^j - F_{t+h|t}^j) = \frac{S_{t+h}^j}{F_{t+h|t}^j} - 1, & \text{if } y_{th}^{jm,TS2} = -1, \end{cases}$$

2.4 | Trading strategies

We use two trading strategies in order to define the three single assets (three assets for each trading strategy) that the investor can select from. Trading strategy 1 (TS1) is based on buying the foreign currency if its price is forecast to rise and selling it when its price is forecast to fall (“buy low, sell high strategy”), and trading strategy 2 (TS2) exploits the forward rate unbiased expectation hypothesis and is related to the so-called carry trade strategy (“carry trade based strategy”).

Trading strategy 1 is a simple “buy low, sell high” trading strategy as described in Gençay (1998), where the selling/buying signal is based on the current exchange rate. Forecast upward movements of the exchange rate with respect to the actual value (positive returns) are executed as long positions, while forecast downward movements (negative returns) are executed as short positions. For each exchange rate model and forecast combination method m and forecast horizon h the trading strategy 1 is defined by the following trading signal, $y_{th}^{jm,TS1}$, and (discrete) return, $r_{t+h,h}^{jm,TS1}$:

$$y_{th}^{jm,TS1} = \begin{cases} 1, & \text{if } \hat{S}_{t+h|t}^{jm} < S_t^j \quad (\text{one FCU of currency } j \text{ is sold at } t \text{ and bought at } t+h) \\ -1, & \text{if } \hat{S}_{t+h|t}^{jm} > S_t^j \quad (\text{one FCU of currency } j \text{ is bought at } t \text{ and sold at } t+h), \end{cases} \quad (8)$$

$$r_{t+h,h}^{jm,TS1} = \begin{cases} \frac{1}{S_t^j} (S_t^j - S_{t+h}^j) = 1 - \frac{S_{t+h}^j}{S_t^j}, & \text{if } y_{th}^{jm,TS1} = 1 \\ \frac{1}{S_t^j} (S_{t+h}^j - S_t^j) = \frac{S_{t+h}^j}{S_t^j} - 1, & \text{if } y_{th}^{jm,TS1} = -1, \end{cases}$$

Trading strategy 2 is based on exploiting the forward rate unbiased expectation hypothesis. In perfect markets, the forward exchange rate is an unbiased predictor of the corresponding future spot exchange rate. If this hypothesis does not hold, a trading strategy based on exchange rate forecasts may earn positive trading profits. This trading strategy thus depends on whether the exchange rate forecast is above or below the forward rate. The trading signal, $y_{th}^{jm,TS2}$, and return, $r_{t+h,h}^{jm,TS2}$, are defined as follows:

where $F_{t+h|t}^j$ is the forward exchange rate (EUR/FCU) at time t with respect to currency j , maturing at time $t+h$.

2.5 | Optimal portfolios

In order to assess whether exchange rate forecasts based on macroeconomic fundamentals improve the profitability of currency portfolios, we investigate the performance of (optimal) currency portfolios of returns implied by two strategies described above, which exploit the potential predictability of exchange rate changes. The optimal portfolio consists of returns implied by a certain trading strategy applied to the three foreign exchange rates EUR/USD, EUR/GBP, and EUR/JPY. We refer to these individual returns as (*single*) *assets* based on the EUR/USD, EUR/GBP, and EUR/JPY exchange rates, respectively. In building optimal portfolios, investors behave according to particular preferences. We model the following types of preferences: mean-variance (MV), conditional value-at-risk (CVaR), linear, linear loss aversion (LLA), and quadratic loss aversion (QLA). As benchmark portfolios, relative to which the optimal portfolios are evaluated,

we consider both (i) the single assets based on individual exchange rates from which the optimal portfolios are composed, as well as (ii) equally weighted (EW) portfolios.⁸

Consider an investor who dynamically (e.g., on a monthly basis) re-balances her portfolio. Let $\mathbf{r}_{th}^{TS} = (r_{th}^{USD,TS}, r_{th}^{GBP,TS}, r_{th}^{JPY,TS})'$, where $r_{th}^{USD,TS}$ is the return at time t implied by trading strategy TS and exchange rate forecasts of the EUR/USD for horizon h . Similarly, $r_{th}^{GBP,TS}$ is the return based on EUR/GBP exchange rate forecasts,

⁸A portfolio with equal weights was also investigated in Burnside et al. (2008) and Burnside, Eichenbaum, and Rebelo (2011), using the carry trade strategy (exploiting the forward rate unbiased expectation hypothesis) and the momentum strategy (stipulating to sell when it was profitable to sell before).

and $r_{th}^{JPY,TS}$ is the return based on EUR/JPY exchange rate forecasts. Let $\mathbf{x}_{th}^{TS} = (x_{th}^{USD,TS}, x_{th}^{GBP,TS}, x_{th}^{JPY,TS})'$, where $x_{th}^{i,TS}$ denotes the proportion of wealth invested at time t in trading strategy TS , whose returns are implied by MSE-based composite forecasts of exchange rate i for horizon h . Returns are constructed with respect to composite forecasts based on the traditional MSE. This means that forecasting models (or forecast combination methods) are chosen such that the MSE is the smallest over a span of 12 months, a period that appears to be appropriate in order to capture changing market conditions.⁹

In our application the returns are available from January 2005 until January 2016 at a monthly frequency, and the optimization exercises are performed for a 3-year rolling window; that is, we optimize over 36 observations. The evaluation period is therefore January 2008 to January 2016 (97 observations). The portfolio optimization exercises are carried out under the following preference schemes:

- *Mean-variance (MV) preferences.* We consider investors that minimize the variance of their portfolio:

$$\min_{\mathbf{x}_{th}^{TS}} \left\{ (\mathbf{x}_{th}^{TS})' \hat{\Sigma}_{t+h|t}^{TS} \mathbf{x}_{th}^{TS} \mid \mathbf{0} \leq \mathbf{x}_{th}^{TS} \leq \mathbf{1}, \mathbf{1}' \mathbf{x}_{th}^{TS} = 1 \right\}, \quad (10)$$

where $\mathbf{0} = (0, 0, 0)'$, $\mathbf{1} = (1, 1, 1)'$, and $\hat{\Sigma}_{t+h|t}^{TS}$ is the estimate of the (3×3) -dimensional conditional covariance matrix of returns implied by the trading strategy TS and model (or forecast combination method or composite forecasts) m .¹⁰

- *Conditional value-at-risk (CVaR) preferences.* We consider an investor that maximizes the conditional expectation of the left tail of portfolio return distribution such that portfolio returns do not exceed the β 's quantile of portfolio return, that is:

$$\max_{(\mathbf{x}_{th}^{TS}, \alpha_\beta)} \left\{ \mathbb{E} \left((\mathbf{x}_{th}^{TS})' \mathbf{r}_{th}^{TS} \right) \mid (\mathbf{x}_{th}^{TS})' \mathbf{r}_{th}^{TS} \leq \alpha_\beta, \right. \\ \left. \mathbf{0} \leq \mathbf{x}_{th}^{TS} \leq \mathbf{1}, \mathbf{1}' \mathbf{x}_{th}^{TS} = 1 \right\}, \quad (11)$$

where $\beta \in (0, 1)$ and α_β is the β -quantile of the portfolio return. Problem 11 is equivalent to¹¹

$$\max_{(\mathbf{x}_{th}^{TS}, \alpha_\beta)} \left\{ \alpha_\beta - \frac{1}{t\beta} \sum_{l=1}^t \left[\alpha_\beta - (\mathbf{x}_{th}^{TS})' \mathbf{r}_{th}^{TS} \right]^+ \right\} \\ \left. \mathbf{0} \leq \mathbf{x}_{th}^{TS} \leq \mathbf{1}, \mathbf{1}' \mathbf{x}_{th}^{TS} = 1 \right\}, \quad (12)$$

⁹In a related analysis using similar returns, we also examined longer and shorter time periods to determine the composite forecast. First we looked at the performance results based on using the total period up to time t , which were usually worse than those based on a period of 12 months. Then we experimented with a period of 6 months, which in most cases resulted in similar or lower performance measures.

¹⁰For a seminal presentation of the mean-variance model, see Markowitz (1952).

¹¹See Rockafellar and Uryasev (2000) for more details.

where $[t]^+$ denotes the maximum of 0 and t . In our application we take $\beta = 0.05$.

- *Linear preferences.* Investors with linear utility functions maximize the expected return of their portfolio:

$$\max_{\mathbf{x}_{th}^{TS}} \left\{ \mathbb{E} \left((\mathbf{x}_{th}^{TS})' \mathbf{r}_{th}^{TS} \right) \mid \mathbf{0} \leq \mathbf{x}_{th}^{TS} \leq \mathbf{1}, \mathbf{1}' \mathbf{x}_{th}^{TS} = 1 \right\}. \quad (13)$$

We denote this investor the “linear” investor.

- *Linear loss aversion (LLA) preferences.* Loss aversion, which is a central finding of prospect theory (see Kahneman & Tversky, 1979) describes the fact that people are more sensitive to losses than to gains, relative to a given reference point \hat{y}_t . The simplest form of such loss aversion is linear loss aversion, where the marginal utility of gains and losses is fixed.¹² Linear loss aversion preferences can be modeled as

$$\max_{\mathbf{x}_{th}^{TS}} \left\{ \mathbb{E} \left((\mathbf{x}_{th}^{TS})' \mathbf{r}_{th}^{TS} - \lambda \left[\hat{y}_t - (\mathbf{x}_{th}^{TS})' \mathbf{r}_{th}^{TS} \right]^+ \right) \mid \right. \\ \left. \mathbf{0} \leq \mathbf{x}_{th}^{TS} \leq \mathbf{1}, \mathbf{1}' \mathbf{x}_{th}^{TS} = 1 \right\}, \quad (14)$$

where $\lambda > 0$ is the loss aversion (or penalty) parameter. Under the given utility, investors face a trade-off between return on the one hand and shortfall below the reference point on the other hand. Interpreted differently, the utility function contains an asymmetric or downside risk measure, where losses are weighted differently from gains. In our application we take the zero return as the reference point; that is, $\hat{y}_t = 0$, and $\lambda = 1.25, 5$.

- *Quadratic loss aversion (QLA) preferences.* Under quadratic loss aversion preferences, large losses are punished more severely than under linear loss aversion preferences.¹³ The quadratic loss aversion preferences can be modeled as

$$\max_{\mathbf{x}_{th}^{TS}} \left\{ \mathbb{E} \left((\mathbf{x}_{th}^{TS})' \mathbf{r}_{th}^{TS} - \lambda \left(\left[\hat{y}_t - (\mathbf{x}_{th}^{TS})' \mathbf{r}_{th}^{TS} \right]^+ \right)^2 \right) \mid \right. \\ \left. \mathbf{0} \leq \mathbf{x}_{th}^{TS} \leq \mathbf{1}, \mathbf{1}' \mathbf{x}_{th}^{TS} = 1 \right\}, \quad (15)$$

where the notation is the same as in the LLA preferences; that is, in our application we take again the zero return as the reference point, $\hat{y}_t = 0$, and $\lambda = 1.25, 5$. The problems given by Equations 14 and 15 are equivalent to higher-dimensional linear programming problems (see Fortin & Hlouskova, 2015).

¹²The optimal asset allocation decision under linear loss aversion has been extensively studied; see, for example, Gomes (2005), He and Zhou (2011), and Fortin and Hlouskova (2011).

¹³The penalty on losses under quadratic loss aversion is also referred to as quadratic shortfall.

2.6 | Performance

2.6.1 | Performance measures

The main performance measures we consider are the mean return, the Omega measure, the Sharpe ratio (SR) and the Sortino ratio (SoR), where the last three measures are risk adjusted and in that sense reflect better the overall considerations of a typical investor (who is not only interested in the portfolio return but also in the portfolio risk). These measures are, among others, reported in our empirical results. The Omega measure is the upside potential of the return with respect to its downside potential relative to the zero return, $\text{Omega} = \sum_{i=1}^n |\max\{r_i, 0\}| / \sum_{i=1}^n |\min\{r_i, 0\}|$, where $i = 1$ corresponds to January 2008 and $i = n$ to January 2016. A larger ratio indicates that the asset provides the chance of more gains relative to losses. The Omega measure is a risk-adjusted performance measure which considers all moments of the return distribution, whereas the Sharpe and Sortino ratios only consider the first two moments (a modified version of the second moment in the case of the Sortino ratio).¹⁴ The Sharpe ratio is the mean divided by the volatility of the return, the Sortino ratio is a modified version of the Sharpe ratio which uses downside volatility with respect to the zero return (instead of standard deviation) as the denominator, i.e., $\text{SR} = \bar{r}/\sigma$ and $\text{SoR} = \bar{r}/\sigma_D$, where \bar{r} and σ are the mean and the standard deviation of r_t calculated with respect to the sample January 2008 through January 2016, and σ_D is the downside volatility (with respect to the zero return) calculated as $\sigma_D = \sqrt{\frac{1}{n} \sum_{i=1}^n (\min\{r_i, 0\})^2}$, where $i = 1$ corresponds to January 2008 and $i = n$ to January 2016.¹⁵ The natural benchmark return for our application appears to be a zero return, reflecting that the investor does not take any position in the foreign exchange market. In both cases a larger ratio indicates a higher return per unit of risk, so usually higher figures denote a better performance. This does not apply, however, when the mean return is negative.

2.6.2 | Bootstrap test for equal performance

In order to assess whether the performance superiority of certain (optimal) portfolios is systematic and not due to luck, we perform the bootstrap stepwise multiple superior

predictive ability test (stepM-SPA) by Hsu, Hsu, and Kuan (2010) for the comparison of optimal portfolio performance with respect to the benchmark models. The test is based on the bootstrap method of Politis and Romano (1994), the stepwise test of multiple check by Romano and Wolf (2005) and the test for superior predictive ability of Hansen (2005).¹⁶

The following relative performance measures, $d_{\text{opt},th}^{\text{TS}}$, $\text{opt} = \text{MV, CVaR, linear, LLA, and QLA}$ with $\lambda = 1.25, 5$, $t = \text{January 2008 to January 2016}$, $h = 1, 3, 6, 12$, are computed and the tests are defined based alternatively on $d_{\text{opt},th}^{\text{TS}} = r_{\text{opt},th}^{\text{TS}} - r_{B,th}^{\text{TS}}$ and $d_{\text{opt},th}^{\text{TS}} = \text{Omega}_{\text{opt},th}^{\text{TS}} - \text{Omega}_{B,th}^{\text{TS}}$.

As benchmark portfolios, for which the corresponding measures (denoted by subindex B) are calculated, we consider the *single assets* based on the returns implied by a certain trading strategy which is applied to single exchange rates, as well as the equally weighted portfolio (EW). The single assets and the equally weighted portfolio are either based on composite forecasts or on the random walk. Note that the performance measure related to Omega is defined as $\text{Omega}_t = \max\{r_t, 0\} / (\frac{1}{n} \sum_{i=1}^n |\min\{r_i, 0\}|)$, where we skip indices and superscripts to simplify the notation. Note that the bootstrap test cannot be performed for the Sharpe and the Sortino ratios, as for negative values it is not true that larger values are associated with a better performance (the test involves the ordering of calculated statistics).¹⁷

The bootstrap stepM-SPA test is a comprehensive test across all portfolio optimization models under consideration and directly quantifies the effect of data snooping by testing the null hypothesis that the performance of the best model is no better than the performance of the benchmark model. For the seven optimization models $\text{opt} = \text{MV, CVaR, linear, LLA}$ with $\lambda = 1.25, 5$, and QLA with $\lambda = 1.25, 5$; we consider the test of $H_0^{\text{opt}} : \mathbb{E}(d_{\text{opt},th}^{\text{TS}} \leq 0)$ against $H_A^{\text{opt}} : \mathbb{E}(d_{\text{opt},th}^{\text{TS}} > 0)$. In our empirical application we use the output of the test to identify those optimal portfolios that outperform the benchmark portfolio at a certain significance level.

2.6.3 | Break-even transaction costs

The break-even transaction cost, τ_{be} , is that constant (proportional) transaction cost that makes the net gain (net of transaction cost) of the optimal portfolio equal to the net gain of the benchmark portfolio¹⁸:

¹⁴The formal definition of the Omega measure is $\frac{\int_{tr}^{\infty} (1-F(x))dx}{\int_{-\infty}^{tr} F(x)dx} = \frac{\mathbb{E}(X-tr|X>tr)P(X>tr)}{\mathbb{E}(tr-X|X\leq tr)P(X\leq tr)}$, where $F(\cdot)$ is the cumulative distribution function of returns and tr is the target return, which in our case is zero.

¹⁵In our empirical results we also present a performance measure labeled “downside volatility ratio” that gives the proportion of downside volatility to the downside and upside volatility σ_U , where $\sigma_U = \sqrt{\frac{1}{n} \sum_{i=1}^n (\max\{r_i, 0\})^2}$, and thus the ratio is calculated as $\frac{\sigma_D}{\sigma_D + \sigma_U}$.

¹⁶For more details on the test, see Hsu et al. (2010).

¹⁷For example, with a higher volatility the Sharpe ratio *increases* instead of *decreases*, given that the mean is negative.

¹⁸Break-even transaction costs are also considered in Han (2006), Della Corte et al. (2009), and Della Corte and Tsiakas (2013), for example. Our computations are based on the transaction value (over the total period of the investment strategy) whereas their computations are based on (one-period) returns.

$$\tau_{be} = \frac{\sum_{t=2}^T (V_t^{\text{opt}} - V_{t-1}^{\text{opt}}) - \sum_{t=2}^T (V_t^{\text{B}} - V_{t-1}^{\text{B}})}{\sum_{t=2}^T V_t^{\text{opt}} \sum_{i=1}^n \left| x_{i,t}^{\text{opt}} - \frac{x_{i,t-1}^{\text{opt}} (1+r_{i,t}^{\text{opt}})}{1+r_{p,t}^{\text{opt}}} \right| - \sum_{t=2}^T V_t^{\text{B}} \sum_{i=1}^m \left| x_{i,t}^{\text{B}} - \frac{x_{i,t-1}^{\text{B}} (1+r_{i,t}^{\text{B}})}{1+r_{p,t}^{\text{B}}} \right|} \quad (16)$$

where V_t is the value of the (optimal/benchmark) portfolio at time t , $x_{i,t}$ is the weight of asset i at time t in the (optimal/benchmark) portfolio, $r_{i,t}$ is the return of asset i at time t in the (optimal/benchmark) portfolio, and $r_{p,t}$ is the return of the (optimal/benchmark) portfolio at time t . Superscripts opt and B denote the optimal and the benchmark portfolios, respectively. We assume that the initial investment in both portfolios is the same; that is, $V_1^{\text{opt}} = V_1^{\text{B}}$.¹⁹ We compare all optimal portfolios (MV, CVaR, linear, $\text{LLA}_{\lambda=1.25}$, $\text{LLA}_{\lambda=5}$, $\text{QLA}_{\lambda=1.25}$, and $\text{QLA}_{\lambda=5}$) with the three single assets and the EW portfolio, based on composite forecasts and on the random walk, respectively.

If the actual transaction cost is lower than the break-even transaction cost then the optimal portfolio achieves a larger gain than the benchmark portfolio, after controlling for transaction costs.²⁰ So larger break-even transaction costs indicate that the optimal portfolio outperforms the benchmark portfolio more easily, while negative break-even transaction costs suggest that the optimal portfolio does not outperform the benchmark portfolio. For the discussion of our empirical results we use a value of 0.1% for the actual transaction cost.²¹

3 | EMPIRICAL RESULTS

3.1 | Data, estimation, and predictions

We base our empirical analysis on monthly data spanning the period from January 1980 until January 2016 for the EUR/USD, EUR/GBP, and EUR/JPY exchange rates. The beginning of the sample is thus $T_0 = \text{January 1980}$; the beginning of the hold-out forecasting sample for individual models used in order to obtain weights based on predictive accuracy is given by $T_1 = \text{January 2000}$. The

beginning of the actual out-of-sample forecasting sample is $T_2 = \text{January 2005}$, and the end of the data sample is $T_3 = \text{January 2016}$. All data are obtained from Thomson Reuters Datastream.²²

Given the large set of statistics computed for the analysis, we start by describing how the results are depicted in the tables. We first report results on the performance of optimal currency portfolios constructed by different types of investors, namely MV, CVaR, linear, linear loss averse ($\text{LLA}_{\lambda=1.25}$, $\text{LLA}_{\lambda=5}$), and quadratic loss averse ($\text{QLA}_{\lambda=1.25}$, $\text{QLA}_{\lambda=5}$) investors (the first seven data columns in the tables of results). The optimal portfolios consist of three single assets, the returns of which are implied by a certain trading strategy (from the two trading strategies described above) with respect to MSE-based composite forecasts of the EUR against the USD, GBP, and JPY. In addition, we present results on how these optimal portfolios compare to *benchmark portfolios*, which include the three single assets that compose the portfolio as well as the equally weighted portfolio (the following four data columns in the tables of results). We also consider as benchmark portfolios the three single assets which are based on the simple random walk (RW) model to forecast exchange rates (instead of on composite forecasts) as well as the equally weighted portfolio of these three assets (the last four data columns).

The tables are structured in five horizontal blocks. The first block presents the four main performance measures on which we base our analysis: the mean return, the Omega measure, the Sharpe ratio, and the Sortino ratio. For two of them, the mean and the Omega measure, we also performed the bootstrap-based stepM-SPA test of Hsu et al. (2010). The second block shows additional statistical descriptives, namely median, volatility, downside volatility, downside volatility ratio, CVaR, skewness, and kurtosis. The third block presents break-even transaction costs for the optimal portfolios, first in relation to the benchmark portfolios based on composite forecasts (where the evaluation criterion is the MSE over the last 12 months) and then in relation to the benchmark portfolios based on the random walk. For the discussion of our empirical results we use a value of 0.1% for the actual transaction cost. The fourth block shows the realized returns over the last 5, 3, and 1 years, and the last block gives the mean portfolio allocations.

¹⁹The break-even transaction cost is not sensitive with respect to the starting value.

²⁰In fact this is only true if the trading cost for the optimal portfolio is larger than the trading cost for the benchmark portfolio (i.e., if the denominator in Equation 16 is positive), which is always true in our applications. Note that it is often a typical feature of the benchmark portfolio that the involved amount of trading is rather small.

²¹Marquering and Verbeek (2004), for example, consider three levels of transaction costs, 0.1%, 0.5%, and 1%, to represent low, medium, and high costs. They consider equity and bond transactions, however, while we consider currency transactions. The latter usually involve lower transaction costs. For example, the average bid-ask spread related to the middle price over our sample period ranges from 0.02% for the EUR/USD exchange rate to 0.06% for the EUR/JPY exchange rate. So a transaction cost of 0.1% seems to be a quite high level in the foreign exchange market.

²²Details on the sources for all variables used are given in the Appendix.

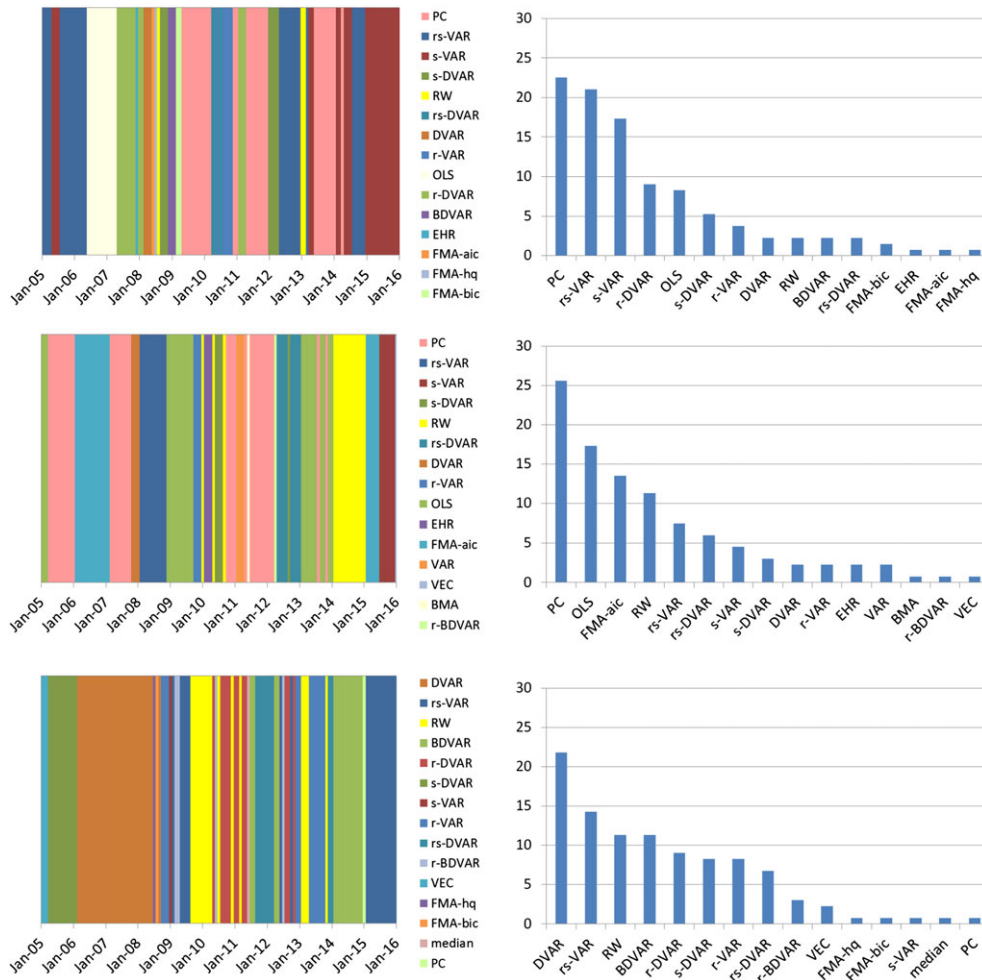


FIGURE 1 MSE-based best forecast models for 1-month-ahead forecasts. The figure shows the best forecast models, that is, the models generating the composite forecasts (left) and the distribution of the best models, that is, the number of times different forecast models are selected over the total period, in percent (right). From top to bottom the figure presents best models for the EUR/USD, the EUR/GBP, and the EUR/JPY [Colour figure can be viewed at wileyonlinelibrary.com]

The subindices in the first two rows of the first block of the benchmark portfolios show the results of the bootstrap test. Their values indicate how many optimal portfolios outperform (in terms of the mean return and the Omega measure) that specific benchmark portfolio. If no subindex is present, no null hypothesis is rejected; that is, the benchmark portfolio is not outperformed by any of the particular optimal portfolios used. If there is only one subindex, its value indicates the number of optimal portfolios that outperform the benchmark portfolio at the 10% significance level. In the case of two subindices, the first one indicates the number of optimal portfolios that outperform the benchmark portfolio at the 5% significance level and the second one at the 10% significance level. In general, statements relating to the bootstrap test will be made using the 10% significance level.

The results are reported for forecast horizons of 1 and 3 months, for both trading strategies, and for composite exchange rate forecasts based on the minimum mean

square error over the last twelve months. Additional results for forecast horizons of 6 and 12 months are not discussed in detail but we point out the main differences with respect to shorter forecast horizons, with the corresponding tables presented in the supporting information Appendix. To allow for a better reading flow, we sometimes simply state USD when we mean the single asset based on the returns implied by the EUR/USD exchange rate (forecasts) for a given trading strategy (analogously for GBP and JPY). To simplify the notation, we sometimes skip the word “portfolio” and just refer to MV instead of MV portfolio (analogously for CVaR, linear, LLA, and QLA).

3.2 | A snapshot example: From composite forecasts to single asset returns to the optimal portfolio

To get a better picture of the process leading up to the optimal portfolio choice, we describe the required steps for a concrete example, where we consider a forecast horizon of

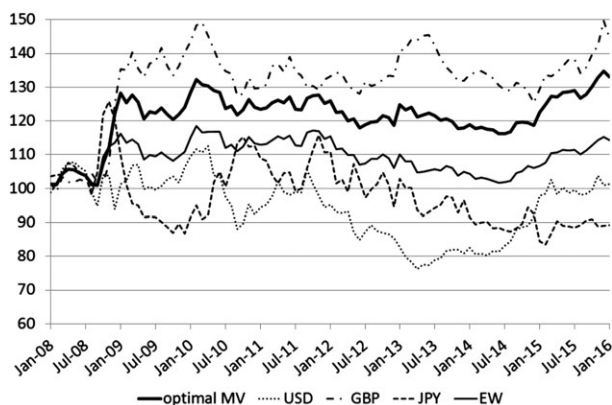


FIGURE 2 Optimal MV portfolio, single assets and EW portfolio based on a 1-month forecast horizon and TS1. The figure shows indices (with a starting value of 100) based on the returns of the optimal mean–variance portfolio, the single assets, and the equally weighted portfolio for trading strategy 1 and a forecast horizon of 1 month

1 month and a mean–variance investor. We use MV investing because it is a popular way of selecting assets and, in addition, the MV portfolio often achieves the highest performance measures. We start by presenting the models and forecast combination methods that are chosen as the best ones, based on the minimum MSE over the last 12 months, for each exchange rate (MSE-based composite forecasts). We then show the returns implied by these composite forecasts and trading strategy 1, for the single assets, the equally weighted portfolio, and the optimal portfolio. Finally we look at the resulting optimal portfolio weights.

The left part of Figure 1 shows how the best forecast models change over the period January 2005 to January 2016, while the right part presents aggregated information on the number of times (percent) a given forecast model is chosen over time. From top to bottom we present the results based on the EUR/USD, the EUR/GBP, and the EUR/JPY exchange rates. For each case, the list of best models includes 15 forecast models/forecast combinations from the total list of 25, but not all these 15 models coincide across currencies. Even though it is visible how models change, and sometimes very quickly (although most of the time best models remain the best for a while), some appear to be chosen quite frequently. The principal components (PC) method, for example, is the model most often selected as the best one for the USD (23%) and for the GBP (26%). On the other hand, it appears only once as the best model for the JPY.²³ The random walk is chosen as the best forecast model only 2% of the time for the USD, while it is selected as the best model 11% of the time for both the GBP and the JPY.

If one is interested in how often single models are chosen as opposed to forecast combination methods, rather surprisingly, single models appear particularly powerful in forecasting. For the USD, single models are chosen 65% of the time and for the JPY even 97% of the time. For the GBP, on the other hand, single models are selected less often, namely 41% of the time. The results for the USD and the JPY (approximately) carry over to a forecast horizon of 3 months, while the results for the GBP are reversed. So for a forecast horizon of 3 months, individual models beat the forecast combination methods in all three currencies. This is also true for longer forecast horizons.²⁴

Figure 2 plots the indices corresponding to the cumulative returns of the optimal MV portfolio together with those of the three single assets that make up the portfolio and the equally weighted portfolio for trading strategy 1 and a forecast horizon of 1 month. The starting values of the indices are set to 100. The graph shows that over the period ranging from January 2008 to January 2016 the total return is highest for the single asset based on the GBP, followed by the optimal MV portfolio and the equally weighted portfolio. The USD shows a return which is barely positive and the JPY shows a clearly negative return. The mean returns are in line with the total return one can see in the graph. The GBP achieves the highest mean return (5.2%), followed by the MV optimal portfolio (3.9%) and the equally weighted portfolio (1.9%). If one takes into account risk considerations and hence looks at, for example, the Sharpe ratio, then the MV portfolio outperforms all other assets as well as the equally weighted portfolio. This is due to the volatility, which is clearly larger for the GBP than for the optimal portfolio. Descriptive statistics relating to the returns shown in the graph, including the mean, the volatility and the Sharpe ratio, are presented in Table 2. The visual impression that the optimal portfolio beats the single assets based on the USD and the JPY is confirmed by the result of the bootstrap test, which concludes that the MV optimal portfolio outperforms both the USD and the JPY at the 10% significance level in terms of the Omega measure. The graph shows that the total return of the MV portfolio exceeds that of the equally weighted portfolio, that of the USD, and that of the JPY. Taking transaction costs into account, this is still true (provided that a transaction cost of 0.1%, or 1%, is assumed).

Figure 3 depicts the optimal weights of the single assets based on the three exchange rates for the mean–variance (MV) and the conditional value-at-risk (CVaR) investors. We still consider trading strategy 1 and a 1-month forecast

²³For a forecast horizon of 3 months, however, the PC is picked up more often as the best forecast model for the JPY, namely 5% of the time.

²⁴More precisely, individual models are chosen as best models 63% (56%) of the time for the USD, 64% (59%) for the GBP, and 89% (89%) of the time for the JPY, for a forecast horizon of 6 (12) months.

TABLE 2 Optimal currency portfolios: out-of-sample evaluation and comparison with benchmark portfolios (TSI, $h = 1$).

	MV	CVaR, β 0.05	Linear	LLA, λ		QLA, λ		USD	GBP	JPY	EW	USD RW	GBP RW	JPY RW	EW RW
				1.25	5.00	1.25	5.00								
<i>Performance measures</i>															
Mean	3.86	3.66	-1.16	0.11	1.45	-0.86	-0.92	0.84	5.16	-0.36	1.86	-2.75	-0.75	0.41	-1.04 ₂
Omega	1.55	1.54	0.93	1.01	1.17	0.95	0.94	1.06 _{2,2}	1.53	0.98 _{1,2}	1.27	0.83 _{2,3}	0.94 ₂	1.02 ₂	0.89 _{2,2}
Sharpe ratio	0.15	0.16	-0.03	0.00	0.06	-0.02	-0.02	0.02	0.14	-0.01	0.08	-0.07	-0.02	0.01	-0.05
Sortino ratio	0.30	0.29	-0.04	0.01	0.10	-0.03	-0.04	0.03	0.30	-0.01	0.14	-0.09	-0.03	0.01	-0.07
<i>Additional descriptive statistics</i>															
Median	2.02	2.48	-1.13	0.04	0.31	-1.13	-2.01	4.65	1.27	0.78	1.85	-2.32	-4.66	2.48	-1.29
Volatility	7.08	6.55	12.53	8.80	6.95	12.34	11.18	11.67	10.08	14.60	6.29	11.64	10.19	14.60	6.56
Down. vol.	3.67	3.58	8.14	5.36	4.34	7.85	7.02	8.29	4.84	9.64	3.90	8.67	6.21	9.82	4.60
Down. vol. ratio	0.37	0.39	0.46	0.44	0.45	0.45	0.45	0.50	0.35	0.47	0.44	0.53	0.44	0.48	0.50
CVaR, $\beta = 0.05$	-32.73	-32.99	-60.49	-43.41	-35.81	-57.31	-50.88	-62.63	-37.79	-65.27	-35.69	-64.02	-50.12	-65.28	-36.92
Skewness	1.18	0.79	0.98	1.71	0.48	1.10	1.36	-0.13	2.08	0.63	0.80	-0.05	1.86	0.49	0.48
Kurtosis	6.11	4.61	6.69	11.49	3.39	6.84	8.84	3.72	12.24	4.58	6.84	3.70	13.19	4.54	4.49
<i>Break-even transaction costs</i>															
USD	4.52	4.49	-1.56	-0.08	0.74	-1.24	-0.78								
GBP	-1.98	-2.13	-6.42	-2.72	-2.56	-5.78	-3.93								
JPY	6.09	6.09	-0.39	0.56	1.53	-0.15	-0.01								
EW	4.05	4.00	-4.10	-1.01	-0.29	-3.45	-2.07								
<i>Break-even transaction costs (RW)</i>															
USD	7.93	7.96	0.98	1.31	2.47	1.13	0.88								
GBP	5.48	5.47	-0.84	0.31	1.23	-0.57	-0.31								
JPY	5.27	5.26	-1.00	0.23	1.12	-0.72	-0.41								
EW	8.63	8.70	-0.82	0.50	1.66	-0.47	-0.19								
<i>Realized return</i>															
Last 5 years	1.51	1.67	-4.08	0.45	0.89	-3.56	-2.90	1.42	2.17	-3.89	0.23	0.45	-4.95	-3.29	-2.28
Last 3 years	2.16	2.49	-5.20	0.42	1.33	-4.94	-4.10	6.94	0.88	-4.60	1.24	0.04	-3.25	-0.63	-1.05
Last year	8.59	8.91	3.50	7.49	8.10	3.50	1.87	3.50	11.70	5.68	7.17	3.50	-2.25	2.87	1.52
<i>Mean allocation</i>															
USD	27.30	18.42	13.40	11.86	18.53	13.40	11.08	100	0	0	33.33	100	0	0	33.33
GBP	48.52	57.86	56.70	60.61	56.58	57.70	64.19	0	100	0	33.33	0	100	0	33.33
JPY	24.19	23.72	29.90	27.53	24.89	28.90	24.73	0	0	100	33.33	0	0	100	33.33

Note. The table reports annual statistics of a monthly reallocated optimal currency portfolio and mean optimal weights, based on an optimization period of 36 months (rolling window), trading strategy 1 and a 1-month forecast horizon. The evaluation period covers January 2008 to January 2016. Statistics are calculated on the basis of monthly returns and then annualized assuming discrete compounding. The same statistics are reported for the benchmark portfolios based on composite forecasts (i.e., the single assets of which the portfolios are constructed and the equally weighted portfolio) and for the benchmark portfolios based on the random walk. The subindices show the results of the bootstrap test. Their values indicate how many optimal portfolios outperform (in terms of the respective performance measure) that specific benchmark portfolio. If no subindex is present, the benchmark portfolio is not outperformed by any of the optimal portfolios. If there is only one subindex, its value indicates the number of optimal portfolios outperforming the benchmark portfolio at the 10% significance level. In the case of two subindices, the first one indicates the number of optimal portfolios outperforming the benchmark portfolio at the 5% significance level and the second one at the 10% significance level. Returns, weights, and transaction costs are given in percent.

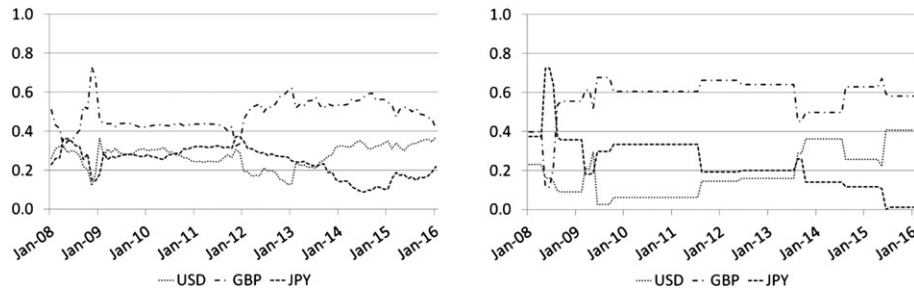


FIGURE 3 Optimal portfolio weights based on a 1-month forecast horizon and TS1. The figure shows the optimal weights of the three single assets in the mean–variance optimal portfolio (left) and in the conditional value-at-risk optimal portfolio (right) for trading strategy 1 and a forecast horizon of 1 month

horizon. On average, the optimal weights in the MV portfolio are 27.3%, 48.5%, and 24.2% for the USD, the GBP, and the JPY. The CVaR portfolio puts even more weight to the GBP (57.9%) on average, thereby reducing the optimal weight in the USD (18.4%), while the optimal weight in the JPY (23.7%) remains basically the same. The two graphs show how the optimal weights evolve over time. We can see that the MV optimal weights change more smoothly over time than the CVaR optimal weights, which remain rather constant over relatively long periods of time and then jump to new weights.

3.3 | Results for the forecast horizon of 1 month

The results relating to a forecast horizon of 1 month are presented in Table 2 for trading strategy 1 and in Table 3 for trading strategy 2.

3.3.1 | Single assets and portfolio composition

The best-performing single asset based on composite forecasts is the GBP, for both trading strategies. The worst performing single asset, for both trading strategies, is the JPY. Even though the ordering of the single assets is the same under TS1 and TS2, the mean returns (and other performance measures) are quite different across the trading strategies. The GBP, for example, is extremely strong under both trading strategies and shows a mean return of 5.2% (TS1) and 6.1% (TS2), respectively. The JPY, on the other hand, shows only a slightly negative mean return (−0.4%) under TS1, while it performs considerably worse under TS2 (−3.6%).²⁵ Note that the volatilities of the single assets are not in line with what one would expect according to the means: The largest mean is actually coupled with the smallest volatility and vice versa. The described proper-

ties of the GBP and the JPY are also reflected in the seven optimal portfolios: Under each trading strategy, the GBP has a large weight in all optimal portfolios (approximately 50% or well above), whereas the JPY shows a considerably smaller weight.

The ordering of single assets based on the random walk is different. In particular, best (worst) assets are not the same any more across the two trading strategies. The best-performing single asset based on the random walk is the JPY under TS1, and the USD under TS2. The worst-performing single asset, on the other hand, is the USD under TS1 and the JPY under TS2. The JPY changes from the best-performing single asset under TS1 to the worst-performing single asset under TS2.

Note that the USD based on the random walk outperforms the USD based on the composite forecast under TS2 (in terms of the mean return, the Omega measure, and the Sharpe and Sortino ratios), where the difference in profitability is quite substantial. Also the JPY based on the random walk outperforms its counterpart based on the composite forecast.²⁶ In all other cases the assets based on the random walk perform worse than their counterparts based on the composite forecast (GBP under both trading strategies, USD under TS1, JPY under TS2). So in most cases there is a positive economic value of forecast models. However, the results might indicate that the USD and the JPY are harder to predict than the GBP.

3.3.2 | Bootstrap analysis

Benchmark portfolios based on composite forecasts. Considering trading strategy 1, the CVaR and MV portfolios outperform the single assets based on the USD and the JPY in terms of the Omega measure. Applying trading strategy 2, again the single assets based on the USD and the JPY are outperformed and, in addition, the equally weighted portfolio (in terms of the Omega measure). While

²⁵The corresponding Omega measures are 1.53 (TS1) and 1.65 (TS2) for the best single asset GBP, and 0.98 (TS1) and 0.82 (TS2) for the worst single asset JPY.

²⁶In fact, the JPY changes from the worst-performing asset based on composite forecasts to the best-performing asset based on the random walk.

TABLE 3 Optimal currency portfolios: out-of-sample evaluation and comparison with benchmark portfolios (TS2, $h = 1$).

	MV	CVaR, β	Linear	LLA, λ		QLA, λ		USD	GBP	JPY	EW	USD	GBP	JPY	EW	USD	GBP	JPY	EW
				1.25	5.00	1.25	5.00												
<i>Performance measures</i>																			
Mean	3.96	3.84	3.47	0.67	1.54	3.57	3.20	1.76	6.06	-3.56	1.35	3.93	-0.25	-5.25 ₂	-0.59				
Omega	1.54	1.52	1.31	1.07	1.18	1.33	1.29	1.12 ₂	1.65	0.82 _{5,6}	1.18 ₂	1.29	0.98 _{2,2}	0.75 _{3,7}	0.94 _{2,2}				
Sharpe ratio	0.15	0.15	0.10	0.03	0.06	0.10	0.09	0.04	0.17	-0.07	0.06	0.10	-0.01	-0.11	-0.02				
Sortino ratio	0.26	0.26	0.16	0.04	0.09	0.17	0.15	0.06	0.36	-0.10	0.08	0.16	-0.01	-0.13	-0.03				
<i>Additional descriptive statistics</i>																			
Median	3.21	2.18	4.28	0.95	1.96	4.96	3.56	5.21	2.64	-0.99	3.55	2.50	-0.54	-2.26	2.06				
Volatility	7.44	7.23	10.35	7.68	7.35	10.27	10.21	11.67	10.03	14.52	6.80	11.63	10.18	14.47	8.33				
Down. vol.	4.25	4.22	6.03	5.04	4.82	6.00	5.94	8.17	4.74	10.54	4.66	7.06	6.53	11.50	5.86				
Down. vol. ratio	0.41	0.42	0.42	0.47	0.47	0.42	0.42	0.50	0.35	0.51	0.49	0.43	0.46	0.57	0.50				
CVaR, $\beta = 0.05$	-40.14	-39.44	-50.28	-43.09	-43.84	-50.28	-50.28	-62.70	-37.58	-68.00	-43.29	-56.92	-52.10	-75.28	-51.13				
Skewness	0.76	0.43	1.32	0.49	0.23	1.33	1.40	-0.15	2.06	0.40	0.15	0.41	1.61	-0.51	0.40				
Kurtosis	6.51	4.19	10.33	4.39	4.71	10.61	10.85	3.80	12.38	4.64	6.94	3.63	13.23	4.32	6.78				
<i>Break-even transaction costs</i>																			
USD	3.52	3.83	1.45	-0.26	0.27	1.29	1.04												
GBP	-3.26	-3.59	-2.91	-3.22	-3.38	-2.35	-2.45												
JPY	8.77	9.58	4.83	2.04	3.09	4.11	3.74												
EW	4.91	5.61	1.65	-0.43	0.16	1.41	1.10												
<i>Break-even transaction costs (RW)</i>																			
USD	0.86	0.92	-0.26	-1.42	-1.16	-0.14	-0.33												
GBP	4.97	5.42	2.39	0.38	1.05	2.07	1.78												
JPY	9.94	10.86	5.59	2.55	3.73	4.74	4.34												
EW	7.70	8.75	3.33	0.68	1.56	2.76	2.39												
<i>Realized return</i>																			
Last 5 years	3.80	4.85	3.27	2.68	2.17	3.35	3.21	2.88	4.46	-2.99	1.77	3.06	2.28	-3.77	0.73				
Last 3 years	5.25	5.58	2.69	3.61	3.63	2.83	2.87	11.19	4.67	-4.56	3.85	6.14	4.00	3.22	4.64				
Last year	7.82	7.96	3.41	6.41	7.63	2.55	3.18	4.10	9.50	5.70	6.67	4.10	-1.51	1.55	1.57				
<i>Mean allocation</i>																			
USD	25.06	14.68	11.34	11.35	15.27	10.20	8.73	100	0	0	33.33	100	0	0	33.33				
GBP	51.61	70.74	75.26	70.36	64.03	76.13	77.06	0	100	0	33.33	0	100	0	33.33				
JPY	23.33	14.58	13.40	18.29	20.70	13.67	14.21	0	0	100	33.33	0	0	100	33.33				

Note. The table reports annual statistics of a monthly reallocated optimal currency portfolio and mean optimal weights, based on an optimization period of 36 months (rolling window), trading strategy 2 and a 1-month forecast horizon. The evaluation period covers January 2008 to January 2016. Statistics are calculated on the basis of monthly returns and then annualized assuming discrete compounding. The same statistics are reported for the benchmark portfolios based on composite forecasts (i.e., the single assets of which the portfolios are constructed and the equally weighted portfolio) and for the benchmark portfolios based on the random walk. The subindices show the results of the bootstrap test. Their values indicate how many optimal portfolios outperform (in terms of the respective performance measure) that specific benchmark portfolio. If no subindex is present, the benchmark portfolio is not outperformed by any of the optimal portfolios. If there is only one subindex, its value indicates the number of optimal portfolios outperforming the benchmark portfolio at the 10% significance level. In the case of two subindices, the first one indicates the number of optimal portfolios outperforming the benchmark portfolio at the 5% significance level and the second one at the 10% significance level. Returns, weights, and transaction costs are given in percent.

the USD and the EW portfolio are again outperformed by the MV and CVaR portfolios, the JPY is now outperformed by all optimal portfolios except the $LLA_{\lambda=1.25}$ portfolio. The equally weighted portfolio is only outperformed under TS2. Regarding the performance based on the mean, none of the benchmark portfolios is significantly outperformed, neither for TS1 nor for TS2.

Benchmark portfolios based on the random walk. We now observe more rejections than in the case of benchmark portfolios based on composite forecasts, as we expect. Under TS1 all RW-based benchmark portfolios are outperformed in terms of the Omega measure by the MV and CVaR portfolios. The USD is additionally outperformed by another optimal portfolio. Under TS2 all RW-based benchmark portfolios except the USD are outperformed in terms of the Omega measure by the MV and CVaR portfolios. The JPY is in fact outperformed by all optimal portfolios. In terms of the mean only the equally weighted portfolio is outperformed under TS1, and the JPY is outperformed under TS2 (both by the MV and CVaR portfolios). Note that under both trading strategies the equally weighted portfolio based on the random walk shows a clearly worse performance than the corresponding equally weighted portfolio based on composite forecasts. Again, this indicates a positive economic value of forecast models.

3.3.3 | Optimal portfolios across investors and trading strategies

Comparison across investors. For a given trading strategy, the best performance is basically always achieved by the MV portfolio. There is only a single instance when this is not true, namely for trading strategy 1 when the Sharpe ratio is considered. In this case the CVaR optimal portfolio yields the best performance. The second-best optimal portfolio is always the CVaR portfolio (or the MV portfolio in the only case when the CVaR portfolio performs the best). In fact, the MV and CVaR portfolios show a quite similar performance in terms of all four performance measures, under both trading strategies.

Comparison across trading strategies. For a given type of investor, optimal portfolios based on trading strategy 2 perform better than those based on trading strategy 1, except for MV and CVaR investors. For these types of investors the performance results under both trading strategies are quite similar, but actually trading strategy 1 slightly outperforms trading strategy 2 (when considering the Omega measure and the Sortino ratio).

3.3.4 | Break-even transaction costs

Taking transaction costs into account, the MV and CVaR optimal portfolios outperform all benchmark portfolios

(except the GBP single asset based on composite forecasts) under both trading strategies. In fact, in most cases when the mean returns of optimal portfolios exceed the mean returns of benchmark portfolios, the optimal portfolios still outperform benchmark portfolios after controlling for transaction costs. Under trading strategy 2, for example, all optimal portfolios outperform the equally weighted portfolio, after taking transaction costs into account, except the $LLA_{\lambda=1.25}$ investor whose benchmark is the EW portfolio based on composite forecasts.

3.4 | Results for the forecast horizon of 3 months

The results relating to a forecast horizon of 3 months are presented in Table 4 for trading strategy 1 and in Table 5 for trading strategy 2.

3.4.1 | Single assets and portfolio composition

For both trading strategies, the best-performing single asset based on composite forecasts is the GBP (in terms of the mean return, the Omega measure, the Sortino ratio, and the Sharpe ratio). The worst-performing single asset for TS1 and TS2 is the JPY. These results are the same as those for a forecast horizon of 1 month. Even though the ordering of the three single assets is the same under both trading strategies, the performance (of a given asset) is different across these strategies (as in the case of the 1-month time horizon). For example, the mean of the USD is 0.7% under TS1, while it is basically zero (0.1%) under TS2. Also for the JPY the mean is larger under TS1 (−2.2%) than under TS2 (−7.1%). The GBP shows very similar means under the two trading strategies (2.2% and 2.4%). Again, the volatilities are not in line with what one would expect according to the means: the single asset with the largest mean has the lowest volatility and vice versa.

The ordering of single assets based on the random walk is different, as for the 1-month forecast horizon. The best-performing single asset is the USD under TS1, and the USD and GBP under TS2.²⁷ The worst-performing single asset is the JPY under both trading strategies.

Under trading strategy 1 it is always true that the single assets based on the composite forecast show a better performance than their counterparts based on the random walk (in terms of all four performance measures). Under trading strategy 2 this statement is always true except for the JPY, where the RW-based single asset outperforms the

²⁷The performance of the USD and GBP under TS2 is extremely similar. While the mean is slightly larger for the GBP, the Omega measure is slightly larger for the USD. The Sharpe and Sortino ratios are the same for both assets.

TABLE 4 Optimal currency portfolios: out-of-sample evaluation and comparison with benchmark portfolios (TS1, $h = 3$).

	MV	CVaR, β		Linear		LLA, λ		QLA, λ		USD	GBP	JPY	EW	USD RW	GBP RW	JPY RW	EW RW
		0.05		1.25	5.00	1.25	5.00	1.25	5.00								
<i>Performance measures</i>																	
Mean	1.56	1.35	-1.49	0.49	1.28	-1.13	-0.37	0.71	2.24	2.24	2.24	-2.21	0.23 ₁	-0.14	-0.26	-3.13	-1.18
Omega	1.38	1.28	0.85	1.08	1.25	0.88	0.95	1.09	1.42	0.82 _{1,3}	1.05 _{1,1}	0.98	0.96 _{1,3}	0.98	0.96 _{1,3}	0.75 _{1,3}	0.78 _{1,3}
Sharpe ratio	0.11	0.08	-0.06	0.02	0.07	-0.04	-0.02	0.03	0.12	-0.08	0.02	-0.01	0.02	-0.01	-0.01	-0.11	-0.09
Sortino ratio	0.18	0.15	-0.07	0.04	0.12	-0.06	-0.02	0.05	0.23	-0.10	0.02	-0.01	0.02	-0.01	-0.02	-0.13	-0.11
<i>Additional descriptive statistics</i>																	
Median	0.87	0.65	-1.46	-0.62	-0.23	-1.85	-0.56	0.07	0.58	-0.77	0.60	0.86	-2.94	1.80	-2.94	1.80	-0.55
Volatility	7.13	8.14	13.51	9.95	8.77	12.99	11.39	10.88	8.98	14.68	6.81	10.89	9.05	14.64	9.05	14.64	6.69
Down. vol.	4.24	4.59	10.51	6.77	5.43	9.97	8.38	7.70	4.88	11.65	4.99	8.11	5.32	12.20	5.32	12.20	5.53
Down. vol. ratio	0.42	0.41	0.56	0.48	0.44	0.55	0.52	0.50	0.39	0.57	0.52	0.53	0.42	0.60	0.42	0.60	0.60
CVaR, $\beta = 0.05$	-27.07	-27.03	-60.29	-40.87	-34.65	-59.32	-49.81	-42.12	-26.55	-60.70	-31.21	-44.56	-26.88	-61.79	-26.88	-61.79	-35.40
Skewness	1.11	1.46	-0.69	0.29	0.72	-0.68	-0.54	-0.32	1.37	-0.70	-0.54	-0.46	1.62	-0.90	1.62	-0.90	-1.25
Kurtosis	10.80	9.43	6.91	7.87	9.30	7.46	8.76	3.82	8.16	4.54	7.38	3.76	8.75	4.39	8.75	4.39	6.92
<i>Break-even transaction costs</i>																	
USD	4.29	3.13	-3.18	-0.44	1.63	-2.99	-1.91										
GBP	-2.16	-3.45	-6.14	-3.79	-2.66	-6.13	-5.05										
JPY	12.41	11.40	0.53	3.78	7.02	0.97	2.03										
EW	10.66	8.33	-3.89	-0.19	3.14	-3.68	-2.22										
<i>Break-even transaction costs (RW)</i>																	
USD	6.96	5.85	-1.96	0.95	3.40	-1.68	-0.62										
GBP	5.59	4.45	-2.59	0.24	2.49	-2.35	-1.28										
JPY	13.68	12.70	1.11	4.44	7.86	1.59	2.65										
EW	16.59	14.73	-1.67	2.30	6.49	-1.27	0.13										
<i>Realized return</i>																	
Last 5 years	1.81	1.62	-1.80	0.51	1.45	-1.25	-0.80	1.44	2.67	-4.83	0.06	1.83	-2.62	-3.75	-2.62	-3.75	-1.19
Last 3 years	1.83	2.39	-1.84	2.15	2.38	-0.72	0.78	2.42	3.64	-7.53	-0.34	1.93	-0.96	-3.12	-0.96	-3.12	-0.48
Last year	9.32	11.33	12.55	11.23	11.80	12.55	12.14	9.99	12.55	-5.80	5.61	9.99	4.79	-5.34	4.79	-5.34	3.34
<i>Mean allocation</i>																	
USD	29.42	9.17	11.34	21.84	25.27	12.25	13.20	100	0	0	33.33	100	0	33.33	0	100	33.33
GBP	50.72	73.38	63.92	57.14	56.69	65.80	65.73	0	100	0	33.33	0	100	0	100	0	33.33
JPY	19.86	17.45	24.74	21.02	18.04	21.95	21.07	0	0	100	33.33	0	0	100	0	100	33.33

Note. The table reports annual statistics of a monthly reallocated optimal currency portfolio and mean optimal weights, based on an optimization period of 36 months (rolling window), trading strategy 1 and a 3-month forecast horizon. The evaluation period covers January 2008 to January 2016. Statistics are calculated on the basis of monthly returns and then annualized assuming discrete compounding. The same statistics are reported for the benchmark portfolios based on composite forecasts (i.e., the single assets of which the portfolios are constructed and the equally weighted portfolio) and for the benchmark portfolios based on the random walk. The subindices show the results of the bootstrap test. Their values indicate how many optimal portfolios outperform (in terms of the respective performance measure) that specific benchmark portfolio. If no subindex is present, the benchmark portfolio is not outperformed by any of the optimal portfolios. If there is only one subindex, its value indicates the number of optimal portfolios outperforming the benchmark portfolio at the 10% significance level. In the case of two subindices, the first one indicates the number of optimal portfolios outperforming the benchmark portfolio at the 5% significance level and the second one at the 10% significance level. Returns, weights, and transaction costs are given in percent.

TABLE 5 Optimal currency portfolios: out-of-sample evaluation and comparison with benchmark portfolios (TS2, $h = 3$).

	MV	CVaR, β 0.05	Linear	LLA, λ		QLA, λ		USD	GBP	JPY	EW	USD RW	GBP RW	JPY RW	EW RW
				1.25	5.00	1.25	5.00								
<i>Performance measures</i>															
Mean	0.85	1.12	-1.95	0.56	1.34	-0.94	-0.31	0.06	2.36	-7.06 _{7,7}	-1.61 _{3,3}	-0.78	-0.68	-5.58 _{4,6}	-2.37
Omega	1.18	1.21	0.77	1.08	1.24	0.88	0.96	1.01	1.44	0.51 _{7,7}	0.72 _{3,5}	0.91	0.90	0.59 _{4,6}	0.65 _{4,5}
Sharpe ratio	0.06	0.07	-0.09	0.03	0.07	-0.04	-0.01	0.00	0.13	-0.26	-0.11	-0.04	-0.04	-0.20	-0.16
Sortino ratio	0.09	0.11	-0.11	0.03	0.10	-0.05	-0.02	0.00	0.24	-0.29	-0.14	-0.05	-0.05	-0.23	-0.18
<i>Additional descriptive statistics</i>															
Median	0.83	0.96	-2.50	1.13	1.14	-0.91	-0.91	-0.10	1.40	-7.57	-1.00	-0.60	2.24	-5.11	0.44
Volatility	7.38	8.38	10.89	10.87	9.59	11.11	10.71	10.82	8.90	14.11	7.04	10.79	8.97	14.29	7.61
Down. vol.	4.69	4.94	8.95	8.03	6.41	8.82	7.99	7.77	4.86	12.50	5.59	7.94	7.30	12.17	6.63
Down. vol. ratio	0.45	0.42	0.59	0.53	0.47	0.57	0.53	0.51	0.39	0.63	0.57	0.52	0.59	0.61	0.63
CVaR, $\beta = 0.05$	-29.37	-26.62	-51.60	-49.25	-38.78	-51.60	-47.06	-41.59	-25.56	-60.42	-32.20	-43.73	-42.55	-57.92	-39.40
Skewness	0.87	1.27	-1.43	-0.83	0.00	-1.29	-0.67	-0.25	1.33	-0.57	-0.32	-0.28	-1.45	-0.45	-1.08
Kurtosis	9.21	8.25	9.32	10.32	10.33	9.38	9.16	3.69	8.12	4.19	6.48	3.68	8.42	4.32	4.95
<i>Break-even transaction costs</i>															
USD	3.81	4.93	-5.08	1.02	3.66	-2.77	-1.08								
GBP	-6.46	-5.88	-14.19	-7.60	-3.92	-10.77	-7.63								
JPY	15.43	17.15	5.22	10.77	12.24	6.28	6.33								
EW	14.88	18.29	-2.78	7.22	10.16	0.46	2.05								
<i>Break-even transaction costs (RW)</i>															
USD	6.34	7.59	-2.84	3.14	5.53	-0.80	0.53								
GBP	4.50	5.65	-4.47	1.59	4.17	-2.23	-0.64								
JPY	14.39	16.05	4.29	9.89	11.47	5.46	5.67								
EW	14.12	16.78	-0.63	7.75	10.25	1.82	2.94								
<i>Realized return</i>															
Last 5 years	1.42	1.94	-0.58	1.94	2.42	0.13	0.45	0.84	2.89	-8.31	-1.30	2.71	2.95	-3.96	0.82
Last 3 years	1.67	3.01	3.14	3.53	3.65	3.11	3.46	2.63	3.17	-7.53	-0.45	5.23	4.12	-1.74	2.74
Last year	7.59	9.73	9.17	8.87	9.20	9.17	9.17	10.52	9.17	-5.78	4.73	10.52	5.54	-3.50	4.41
<i>Mean allocation</i>															
USD	32.91	16.07	18.56	19.01	25.59	17.73	16.16	100	0	0	33.33	100	0	0	33.33
GBP	50.54	70.06	67.01	68.97	65.21	69.51	71.52	0	100	0	33.33	0	100	0	33.33
JPY	16.56	13.87	14.43	12.02	9.20	12.76	12.31	0	0	100	33.33	0	0	100	33.33

Note. The table reports annual statistics of a monthly reallocated optimal currency portfolio and mean optimal weights, based on an optimization period of 36 months (rolling window), trading strategy 2 and a 3-month forecast horizon. The evaluation period covers January 2008 to January 2016. Statistics are calculated on the basis of monthly returns and then annualized assuming discrete compounding. The same statistics are reported for the benchmark portfolios based on composite forecasts (i.e., the single assets of which the portfolios are constructed and the equally weighted portfolio) and for the benchmark portfolios based on the random walk. The subindices show the results of the bootstrap test. Their values indicate how many optimal portfolios outperform (in terms of the respective performance measure) that specific benchmark portfolio. If no subindex is present, the benchmark portfolio is not outperformed by any of the optimal portfolios. If there is only one subindex, its value indicates the number of optimal portfolios outperforming the benchmark portfolio at the 10% significance level. In the case of two subindices, the first one indicates the number of optimal portfolios outperforming the benchmark portfolio at the 5% significance level and the second one at the 10% significance level. Returns, weights, and transaction costs are given in percent.

single asset based on the composite forecast. These results may indicate that the JPY is harder to predict than the USD and the GBP.

3.4.2 | Bootstrap analysis

Benchmark portfolios based on composite forecasts. For each trading strategy, the equally weighted portfolio as well as the JPY are outperformed by optimal portfolios (in terms of the Omega measure). Under TS1, however, only the MV portfolio outperforms the EW portfolio, and the MV, the CVaR, and the $LLA_{\lambda=5}$ portfolios outperform the JPY, while under TS2 (nearly) all optimal portfolios outperform the benchmark portfolios.²⁸ These results indicate that investors should in fact engage in active optimal portfolio management rather than just follow the naive equally weighted portfolio approach. Note that for the 1-month ahead forecast horizon and TS1, the equally weighted portfolio was not outperformed, which implies that the gains from building optimal currency portfolios based on exchange rate forecasts (with respect to the naive EW portfolio) appear to be sensitive to the forecast horizon considered.

Benchmark portfolios based on the random walk. Both the EW portfolio and the JPY are outperformed by optimal portfolios, under TS1 and TS2, in terms of the Omega measure. However, while under TS1 only the MV, CVaR, and $LLA_{\lambda=5}$ portfolios outperform the two benchmark portfolios, under TS2 nearly all optimal portfolios (except the linear one) outperform these benchmark portfolios. Further, the GBP is outperformed by the MV, the CVaR, and the $LLA_{\lambda=5}$ under TS1. For the mean, only the JPY under TS2 is outperformed, namely by all optimal portfolios except the linear one. Note that under both trading strategies the RW-based equally weighted portfolio shows a worse performance than the equally weighted portfolio based on composite forecasts (in terms of all four performance measures). Under TS1, the difference is quite substantial whereas under TS2 the difference is rather small. This again confirms the positive economic value of forecast models.

3.4.3 | Optimal portfolios across investors and trading strategies

Comparison across investors. Under TS1 the best performance is achieved by MV and CVaR portfolios (in terms of the mean, the Omega measure, the Sharpe ratio, and the Sortino ratio)—similar to the 1-month forecast time horizon. Under TS2, however, the $LLA_{\lambda=5}$ and CVaR portfolios perform best.

Comparison across trading strategies. For the MV, CVaR and linear investors, trading strategy 1 yields a higher performance than trading strategy 2 (in terms of all four performance measures). For all other types of investors, trading strategy 2 delivers better performance measures.

3.4.4 | Break-even transaction costs

Taking transaction costs into account, the MV and CVaR optimal portfolios outperform all benchmark portfolios (except the GBP single asset based on composite forecasts) under both trading strategies. This is completely similar to the case of the 1-month forecast horizon. Again, in most cases when the mean returns of optimal portfolios exceed the mean returns of benchmark portfolios, the optimal portfolios still outperform benchmark portfolios after controlling for transaction costs. Under trading strategy 2, for example, all optimal portfolios outperform the equally weighted portfolio (based on composite forecasts and based on the random walk), after taking transaction costs into account, except for the linear investor.

3.5 | Main results for the forecast horizons of 6 and 12 months

A crucial observation with respect to the forecast horizon in general is that the performance of the single assets, as well as of the optimal portfolios, seems to decrease with the forecast horizon. For horizons of 6 and 12 months the mean returns (of both benchmark and optimal portfolios) are in fact mostly close to zero or negative. Based on the bootstrap results, the equally weighted portfolio and the JPY (both based on the composite forecasts) are mostly outperformed, for both trading strategies. In most cases, this is true in terms of the Omega measure and the mean.

Regarding the types of investors, we observe that with an increasing forecast horizon the group of investors achieving the best or second-best performance seems to be widening to include (in addition to the MV and CVaR investors) also the linear, $LLA_{\lambda=5}$, and QLA investors. Although for a forecast horizon of 6 months the MV and CVaR portfolios are still the best or second-best performing optimal portfolios in most cases, the linear and QLA investors seem to be taking over for a forecast horizon of 12 months.

Note that the number of cases when the performance of benchmark portfolios (single assets and equally weighted portfolio) based on the random walk exceeds the performance of benchmark portfolios based on the composite forecasts increases with a forecast horizon of 12 months. Although all benchmark portfolios based on the random walk are outperformed by their counterparts based on composite forecasts for the 6-month forecast horizon (under both trading strategies and in terms of all four

²⁸More precisely, all except the linear and the $QLA_{\lambda=1.25}$, for the EW benchmark, and all for the JPY.

performance measures), nearly all benchmark portfolios based on the random walk in fact outperform their counterparts based on composite forecasts for the 12-month forecast horizon.²⁹

Taking transaction costs into account, the performance of optimal portfolios for a forecast horizon of 6 months is rather similar to the case of 1-month and 3-month forecast horizons. For a forecast horizon of 12 months there are clearly fewer cases when optimal portfolios outperform benchmark portfolios. In fact, under both trading strategies, none of the optimal portfolios outperforms the equally weighted portfolio based on the random walk.

3.6 | Summary and stylized facts

The performance of benchmark portfolios (single assets and the equally weighted portfolio) based on composite forecasts is better than the performance of benchmark portfolios based on the random walk, for forecast horizons of 1, 3, and 6 months and both trading strategies, with only very few exceptions.³⁰ Thus the economic value of using exchange rate forecast models rather than the naive random walk forecasts is pronounced for forecast horizons of up to 6 months. This is not really the case for a 12-month forecast horizon.

Investors with MV and CVaR preferences perform better, in terms of the mean return (also adjusted for transaction costs), the Omega measure, the Sortino ratio, and the Sharpe ratio, under both trading strategies, and for forecast horizons of 1, 3 and 6 months, than investors that choose the naive equally weighted portfolio based on composite forecasts. This observation is even more pronounced when the benchmark is the equally weighted portfolio based on the random walk. In this case nearly all optimal portfolios outperform the equally weighted portfolio. This indicates that exchange rate forecasting and portfolio optimization imply a positive economic value for MV and CVaR investors, when the benchmark is the equally weighted portfolio and the forecast horizon does not exceed 6 months.

In most cases when optimal portfolios show larger mean returns than benchmark portfolios without transaction costs, they outperform the benchmark portfolios also after controlling for transaction costs. In addition, if one believes that non-naive exchange rate forecasts provide positive economic value then break-even transaction costs with respect to benchmarks based on the random walk should exceed break-even transaction costs with respect to

benchmarks based on composite forecasts. For a forecast horizon of 12 months, however, this is mostly not true; that is, break-even transaction costs are often larger when calculated with respect to the benchmark portfolios based on composite forecasts than when calculated with respect to the benchmark portfolios based on the naive random walk. This is another example of the positive economic value of exchange rate forecast models (for forecast horizons of up to 6 months).

The difference in performance between the two trading strategies is rather small.³¹ Regarding the average optimal weights, the highest proportion for forecast horizons of 1, 3 and 6 months and for both trading strategies is allocated to the single asset based on the EUR/GBP. This is also the best-performing single asset, the one with the lowest volatility, and the only one which performs systematically better when based on the composite forecasts than when based on the naive random walk forecasts, that is, for all forecast horizons of 1, 3, and 6 months. These results imply a particularly large economic value of applying exchange rate models to the EUR/GBP exchange rate

4 | CONCLUSION

Using a comprehensive set of exchange rate models, we analyze whether investors in the foreign exchange markets, whose goal is to maximize risk-adjusted profits, should engage in active portfolio management rather than just follow the naive equally weighted portfolio approach or simply invest in one of the assets composing the portfolio. These assets are defined through returns implied by a trading strategy based on exchange rate forecasts. In general, we find that the benchmark portfolios based on the random walk, that is, the benchmark portfolios which do not make use of any (nontrivial) exchange rate forecasts, are more often outperformed by optimal portfolios than benchmark portfolios based on composite forecasts up to a forecast horizon of 6 months. This indicates indirectly that the use of exchange rate forecasts adds to improving the (risk-adjusted) profitability of the benchmark portfolios we consider, that is, of the single assets and the equally weighted portfolio.

²⁹This is in contrast to the stylized fact that non-naive exchange rate forecasts beat the random walk, if at all, at longer time horizons.

³⁰These are the JPY under trading strategy 1 for a 1-month forecast horizon, the USD under trading strategy 2 for a 1-month forecast horizon, and the JPY under trading strategy 2 for a 3-month forecast horizon.

³¹Clear exceptions to this rule are the linear and the quadratic loss-averse investors with a forecast horizon of 1 month, where the mean return under trading strategy 2 is by over 4 percentage points larger than the mean return under trading strategy 1. Otherwise, the difference in mean returns across the two strategies is smaller than 0.6 percentage points for a forecast horizon of 1 month and smaller than 0.8 percentage points for a forecast horizon of 3 months. For longer forecast horizons (6 and 12 months) trading strategy 1 uniformly outperforms trading strategy 2 (in terms of all performance measures), but note that in most of these cases the mean returns are close to zero or negative.

Our results suggest that the performance of optimal portfolios and of benchmark portfolios decreases with an increasing forecast horizon. With respect to the optimal weights, the highest proportion in optimal portfolios for forecast horizons of 1, 3, and 6 months is allocated to the single asset based on the EUR/GBP. This is the only asset which performs systematically better when based on composite forecasts than when based on the naive random walk forecasts. These findings might suggest a particularly large economic value of applying exchange rate models (and forecast combination methods) to the EUR/GBP exchange rate. The overall difference between the “buy low, sell high” strategy and the carry trade strategy is rather small.

ACKNOWLEDGEMENTS

The authors would like to thank an anonymous referee for very helpful comments on an earlier draft of the paper. Ines Fortin and Jaroslava Hlouskova gratefully acknowledge financial support from Oesterreichische Nationalbank (Anniversary Fund, Grant No. 16250).

ORCID

Jesus Crespo Cuaresma  <http://orcid.org/0000-0003-3244-6560>

REFERENCES

- Abbate, A., & Marcellino, M. (2018). Point, interval and density forecasts of exchange rates with time-varying parameter models. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 181, 155–179.
- Bacchini, F., Ciammola, A., Iannaccone, R., & Marini, M. (2010). *Combining Forecasts for a Flash Estimate of Euro Area GDP*, (Contributi ISTAT, No. 3). Rome, Italy: National Statistical Institute of Italy.
- Barroso, P., & Santa-Clara, P. (2015). Beyond the carry trade: Optimal currency portfolios. *Journal of Financial and Quantitative Analysis*, 50, 1037–1056.
- Berkowitz, J., & Giorgianni, L. (2001). Long-horizon exchange rate predictability. *Review of Economics and Statistics*, 83, 81–91.
- Boudoukh, J., Richardson, M., & Whitelaw, R. F. (2008). The myth of long-horizon predictability. *Review of Financial Studies*, 21, 1577–1605.
- Brandt, M. W., Santa-Clara, P., & Valkanov, R. (2009). Parametric portfolio policies: Exploiting characteristics in the cross-section of equity returns. *Review of Financial Studies*, 22, 3411–3447.
- Burnside, C., Eichenbaum, M., & Rebelo, S. (2008). Carry trade: The gains of diversification. *Journal of the European Economic Association*, 6, 581–588.
- Burnside, C., Eichenbaum, M., & Rebelo, S. (2011). Carry trade and momentum in currency markets. *Annual Review of Financial Economics*, 3, 511–535.
- Carriero, A., Kapetanios, G., & Marcellino, M. (2009). Forecasting exchange rates with a large Bayesian VAR. *International Journal of Forecasting*, 25, 400–417.
- Cheung, Y. W., Chinn, M. D., & Pascual, A. G. (2005). Empirical exchange rate models of the nineties: Are any fit to survive. *Journal of International Money and Finance*, 24, 1150–1175.
- Chinn, M. D., & Meese, R. A. (1995). Banking on currency forecasts: How predictable is change in money. *Journal of International Economics*, 38, 161–178.
- Claeskens, G., & Hjort, N. L. (2008). *Model Selection and Model Averaging*. Cambridge, UK: Cambridge University Press.
- Costantini, M., Crespo Cuaresma, J., & Houskova, J. (2014). *Can macroeconomists get rich forecasting exchange rates?* (IHS Economic Series Working Paper No. 305). Vienna, Austria: WU Vienna University of Economics and Business.
- Costantini, M., Crespo Cuaresma, J., & Hlouskova, J. (2016). Forecasting errors, directional accuracy and profitability of currency trading: The case of EUR/USD exchange rate. *Journal of Forecasting*, 35, 652–668.
- Costantini, M., & Kunst, R. (2011). Combining forecasts based on multiple encompassing tests in a macroeconomic core system. *Journal of Forecasting*, 30, 579–596.
- Costantini, M., & Pappalardo, C. (2010). A hierarchical procedure for the combination of forecasts. *International Journal of Forecasting*, 26, 725–743.
- Crespo Cuaresma, J. (2007). *Forecasting Euro Exchange Rates: How Much Does Model Averaging Help?* (Working Papers in Economics and Statistics 2007–24). Innsbruck, Austria: University of Innsbruck.
- Crespo Cuaresma, J., Fortin, I., & Hlouskova, J. (2017). *Exchange Rate Forecasting and the Performance of Currency Portfolios*, (IHS Economic Series Working Paper No. 326). Vienna, Austria: WU Vienna University of Economics and Business.
- Crespo Cuaresma, J., & Hlouskova, J. (2005). Beating the random walk in central and eastern Europe. *Journal of Forecasting*, 24, 189–201.
- Della Corte, P., Sarno, L., & Tsiakas, I. (2009). An economic evaluation of empirical exchange rate models. *Review of Financial Studies*, 22, 3491–3530.
- Della Corte, P., & Tsiakas, I. (2013). Statistical and economic methods for evaluating exchange rate predictability. In James, J., Marsh, I., & Sarno, L. (Eds.), *Handbook of Exchange Rates* (pp. 239–283). Hoboken, NJ: Wiley.
- DeMiguel, V., Garlappi, L., & Uppal, R. (2009). Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy. *Review of Financial Studies*, 22, 1915–1953.
- Doan, T., Litterman, R., & Sims, C. (1984). Forecasting and conditional projections using realist prior distributions. *Econometric Reviews*, 3, 1–100.
- Dornbusch, R. (1976). Expectations and exchange rate dynamics. *Journal of Political Economy*, 84, 1161–1176.
- Feldkircher, M. (2012). Forecast combination and Bayesian model averaging: A prior sensitivity analysis. *Journal of Forecasting*, 31, 361–376.
- Fernandez, C., Ley, E., & Steel, M. (2001). Benchmark priors for Bayesian model averaging. *Journal of Econometrics*, 100, 381–427.
- Fortin, I., & Hlouskova, J. (2011). Optimal asset allocation under linear loss aversion. *Journal of Banking and Finance*, 35, 2974–2990.
- Fortin, I., & Hlouskova, J. (2015). Downside loss aversion: Winner or loser? *Mathematical Methods of Operations Research*, 81, 181–233.

- Frenkel, J. A. (1976). A monetary approach to the exchange rate: Doctrinal aspects and empirical evidence. *Scandinavian Journal of Economics*, 78, 200–224.
- Gençay, R. (1998). The predictability of security returns with simple technical trading rules. *Journal of Empirical Finance*, 5, 347–359.
- Gomes, F. J. (2005). Portfolio choice and trading volume with loss-averse investors. *Journal of Business*, 78, 675–706.
- Granger, C. W. J., & Ramanathan, R. (1984). Improved methods of combining forecasts. *Journal of Forecasting*, 3, 197–204.
- Han, Y. (2006). Asset allocation with a high dimensional latent factor stochastic volatility model. *Review of Financial Studies*, 19, 237–271.
- Hansen, P. R. (2005). A test for superior predictive ability. *Journal of Business and Economic Statistics*, 23, 365–380.
- He, X. D., & Zhou, X. Y. (2011). Portfolio choice under cumulative prospect theory: An analytical treatment. *Management Science*, 57, 315–331.
- Hjort, N. L., & Claeskens, G. (2003). Frequentist model average estimators. *Journal of the American Statistical Association*, 98, 879–899.
- Hlouskova, J., & Wagner, M. (2013). The determinants of long-run economic growth: A conceptually and computationally simple approach. *Swiss Journal of Economics and Statistics*, 149, 445–492.
- Hooper, P., & Morton, J. (1982). Fluctuations in the dollar: A model of nominal and real exchange rate determination. *Journal of International Money and Finance*, 1, 39–56.
- Hsu, P., Hsu, Y., & Kuan, C. (2010). Testing the predictive ability of technical analysis using a new stepwise test without data snooping bias. *Journal of Empirical Finance*, 17, 471–484.
- Jacobs, H., Müller, S., & Weber, M. (2014). How should individual investors diversify? An empirical evaluation of alternative asset allocation policies. *Journal of Financial Markets*, 19, 62–85.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 263–292.
- Kilian, L. (1999). Exchange rates and monetary fundamentals: What do we learn from long-horizon regressions. *Journal of Applied Econometrics*, 14, 491–510.
- Litterman, R. B. (1986). Forecasting with Bayesian vector autoregressions: Five years of experience. *Journal of Business and Economic Statistics*, 4, 25–38.
- MacDonald, R., & Taylor, M. P. (1994). The monetary model of the exchange rate: Long-run relationships, short-run dynamics and how to beat a random walk. *Journal of International Money and Finance*, 13, 276–290.
- Mark, N. C. (1995). Exchange rates and fundamentals: Evidence on long-horizon prediction. *American Economic Review*, 85, 201–218.
- Mark, N. C., & Sul, D. (2001). Nominal exchange rates and monetary fundamentals: Evidence from a small post-Bretton Woods panel. *Journal of International Economics*, 53, 29–52.
- Markowitz, H. M. (1952). Portfolio selection. *Journal of Finance*, 7, 77–91.
- Marquering, W., & Verbeek, M. (2004). The economic value of predicting stock index returns and volatility. *Journal of Financial and Quantitative Analysis*, 39, 407–429.
- Meese, R. A., & Rogoff, K. (1983). Empirical exchange rate models of the seventies. Do they fit out of sample. *Journal of International Economics*, 14, 3–24.
- Menkhoff, L., & Taylor, M. P. (2007). The obstinate passion of foreign exchange professionals: Technical analysis. *Journal of Economic Literature*, 45, 936–972.
- Papaioannou, E., Portes, R., & Siourounis, G. (2006). Optimal currency shares in international reserves: The impact of the euro and the prospects for the dollar. *Journal of the Japanese and International Economies*, 20, 508–547.
- Politis, D. N., & Romano, J. P. (1994). The stationary bootstrap. *Journal of the American Statistical Association*, 89, 1303–1313.
- Raftery, A. E., Madigan, D., & Hoeting, J. A. (1997). Bayesian model averaging for linear regression models. *Journal of the American Statistical Association*, 92, 179–191.
- Rockafellar, R. T., & Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk*, 2, 21–41.
- Romano, J. P., & Wolf, M. (2005). Stepwise multiple testing as formalized data snooping. *Econometrica*, 73, 1237–1282.
- Rossi, B. (2013). Exchange rate predictability. *Journal of Economic Literature*, 51, 1063–1119.
- Sala-i-Martin, X., Doppelhofer, G., & Miller, R. I. (2004). Determinants of long-term growth: A Bayesian averaging of classical estimates (BACE) approach. *American Economic Review*, 94, 813–835.
- Stock, J. H., & Watson, M. W. (2004). Combination forecasts of output growth in a seven-country data set. *Journal of Forecasting*, 23, 405–430.

SUPPORTING INFORMATION

Additional Supporting Information may be found online in the supporting information tab for this article.

How to cite this article: Crespo Cuaresma J, Fortin I, Hlouskova J. Exchange rate forecasting and the performance of currency portfolios. *Journal of Forecasting*. 2018;37:519–540. <https://doi.org/10.1002/for.2518>

APPENDIX : DATA DESCRIPTION AND SOURCES

All time series have monthly periodicity (January 1980 to January 2016) and have been extracted from Thomson Reuters Datastream. The spot (forward) exchange rates are mid exchange rates of the WM/Reuters closing spot (forward) rates. The variables used for the Euro area, Japan, UK, and USA are as follows:

- Money supply: M1 aggregate, indexed 1990:1 = 100. Seasonally unadjusted.
- Output: Industrial production index 1990:1 = 100.
- Short-term interest rate: 3-month interbank offered rate.
- Long-term interest rate: 10-year government bond yield.

The values of the variables used for the euro area before the introduction of the euro are as calculated and provided by Thomson Reuters Datastream.