

Speculators' Positions and Exchange Rate Forecasts: Beating Random Walk Models*

Young Ju Kim[†] Zhipeng Liao[‡] Aaron Tornell[§]

This Version: May, 2014

Abstract

Speculators' positions in futures markets contain useful information to forecast exchange rates. We extract such information by fitting a microfounded regime switching model, and forecast whether speculators will be increasing or decreasing their positions. We use this predicted state to form both directional and point exchange rate forecasts for the six most traded currency pairs. Over forecasting horizons from 6 to 12 months, our directional forecasts have a 58 percent average success ratio and most of our point forecasts are more accurate than those implied by the random walk models. Forecasting evaluation tests show that our empirical findings are significant.

*We thank Graham Elliott, Olivier Gossner, Pierre Olivier Gourinchas, Jin Hahn, Rosa Matskin, Romain Ranciere, Martin Schneider, Allan Timmermann, and Pierre Olivier Weil for helpful comments.

[†]Department of Economics, UCLA, Mail Stop: 147703, Los Angeles, CA 90095. Email: econoky@ucla.edu.

[‡]Department of Economics, UCLA, 8379 Bunche Hall, Mail Stop: 147703, Los Angeles, CA 90095. Email: zhipeng.liao@econ.ucla.edu.

[§]Department of Economics, UC Los Angeles, 8283 Bunche Hall, Mail Stop: 147703, Los Angeles, CA 90095. Email: tornell@ucla.edu.

Speculators' positions in the futures markets contain useful information to forecast exchange rates over horizons of up to a year. We extract such information by fitting a microfounded autoregressive Markov regime switching model to the speculators' net positions data in the Commitments-of-Traders (COT) report.¹ Using the model we forecast whether speculators will be increasing or decreasing their positions. We then use this predicted state to form both *directional* and *point* exchange rate forecasts. When our model detects that speculators in a given currency—say the Yen—start to increase their net Yen positions, they tend to remain in such an accumulation state for several months, and during this period the Yen tends to appreciate. Analogously, when speculators in the Yen shift to a decumulation state, the Yen tends to exhibit a persistent depreciation.

Our forecasting method combines elements of Engel and Hamilton's (1990) finding of long-swings in the Dollar with elements of Evans and Lyons (2002) finding that the *order-flow* of a group of dealers in the interbank foreign exchange market, forecasts exchange rates over short horizons. While order-flow is typically private information, the COT data we use is publicly available, and has forecasting power over longer horizons. An attractive aspect of the COT report, which we exploit, is that it breaks down the positions in the futures markets into those held by *hedgers* and those held by *speculators*.

Over horizons ranging from 1 to 12 months, our *directional forecasts* have a 58% average success ratio across the six most traded currency pairs vis-a-vis the US Dollar.² The forecast success ratios are particularly accurate at the 6-to-12 months horizons: Excluding the Swiss Franc, they range from 58% for the British Pound to 78% for the Australian Dollar.³

Do our directional forecasts succeed in predicting big moves in exchange

¹The COT report is published weekly by the Commodities and Futures Trade Commission (CFTC).

²We consider the Euro, Japanese Yen, British Pound, Australian Dollar, Canadian Dollar and Swiss Franc.

³Our out-of-sample forecasts start in 1994 and end in February 2013. The starting dates are: 01/20/1995 for the Australian Dollar and British Pound; 09/01/1994 for the Canadian Dollar; 04/15/1994 for the Swiss Franc and Japanese Yen; and 05/04/2001 for the Euro.

rates? To answer this question, we propose a novel test that weights each directional forecast by the realized exchange rate change, and evaluates whether our directional forecasts outperform the driftless random walk forecasts. This test may be more relevant than the standard binomial test to situations faced by investors, who look at the profitability associated with their forecasts. At the 6-month (6m), 9-month (9m) and 12-month (12m) horizons, controlling for auto-correlation using the long-run variance (LRV) estimators proposed in Newey and West (1987) and Andrews (1991), the weighted directional test rejects the random walk null in favour of our model across all currencies, except the Swiss Franc. At the 1-month (1m) and 3-month (3m) horizons, the null is rejected in only 5 out of 12 currency-horizon pairs.⁴ The standard binomial test generates similar results.

To evaluate the difference between our *point-forecasts* and random-walk forecasts, we conduct the Diebold-Mariano-West (DMW) test and the Clark-West (CW) test. Controlling for auto-correlation using the Newey-West LRV estimator, the DMW-test rejects the null over the *6m*, *9m* and *12m* forecasting horizons across all currencies, except for the British Pound at the *12m* horizon and the Swiss Franc at all horizons. The CW-test is even more favorable to our model: over the *6m*, *9m* and *12m* forecasting horizons, it rejects the null across all currencies, except the Swiss Franc. The test results using the Andrews LRV estimator are similar to those based on the Newey-West LRV estimator.

The rest of the paper is organized as follows. Section 1 presents a birds-eye view of our forecasting method. Section 2 presents a literature review. In Section 3, we derive the estimation equation from a theoretical model in which speculators observe a noisy signal of the process that drives exchange rates. By inverting the speculators demand for foreign currency, we obtain the estimation model that recovers the information about the exchange rate trend from the observed data on the speculators' positions. Section 4 esti-

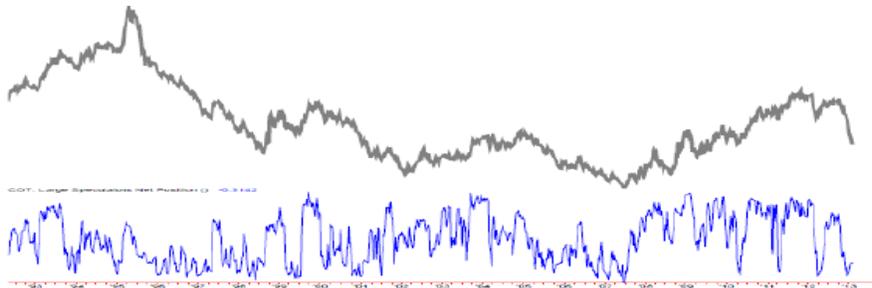
⁴Throughout this paper, when citing the Newey-West LRV estimator, we refer to the kernel smoothed LRV estimator proposed in Newey and West (1987) and bandwidth selection rule suggested in Newey and West (1994). When citing the Andrews LRV estimator, we refer to the kernel smoothed LRV estimator using Bartlett kernel and bandwidth selection rule suggested in Andrews (1991).

mates this autoregressive Markov switching model, using data from the COT report. Section 5 generates the directional exchange rate forecasts and tests two hypotheses to evaluate their performance. Section 6 generates the point forecasts and conducts the DMW-test and the CW-test. Section 7 concludes. Proofs, technical results and the forecasting comparison of our model against the random walk with drift model are given in the appendix for online publication.

1 Outline

Here, we illustrate the gist of our methodology through a series of figures that depict the Yen-Dollar exchange rate together with the forecasting variables we use. This section can be skipped without loss of continuity.

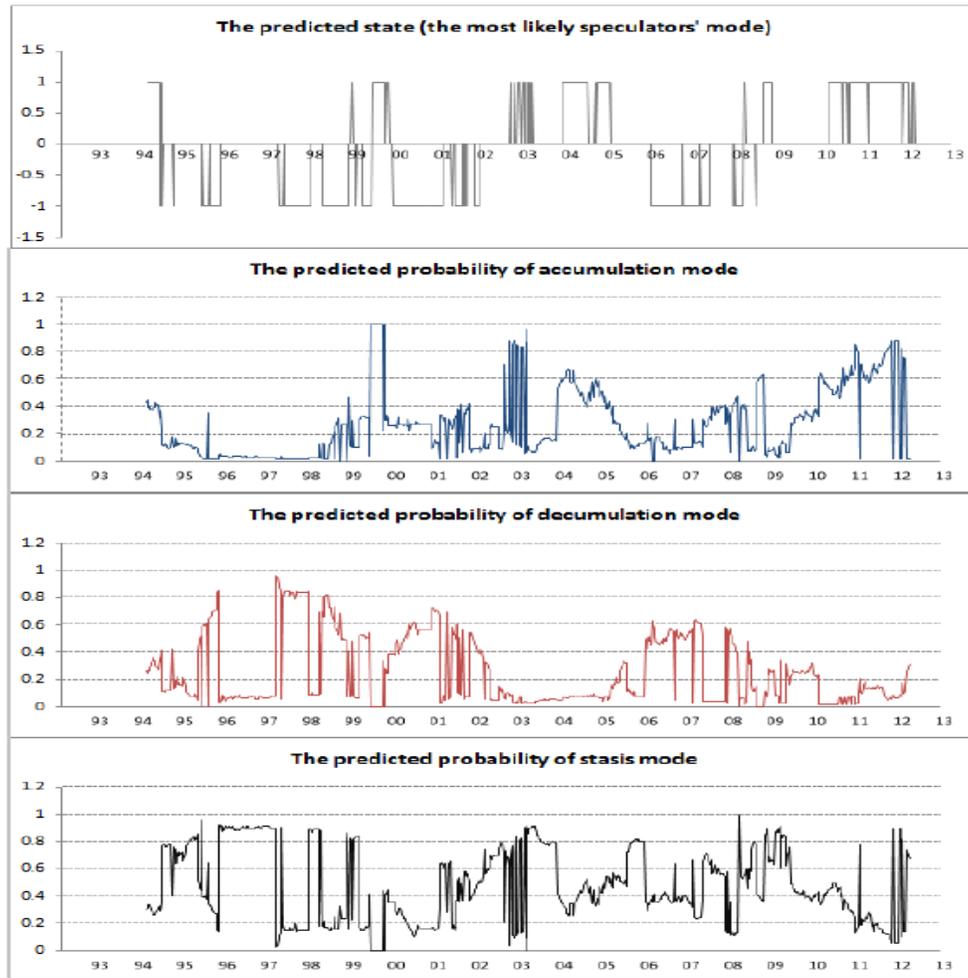
Figure 1. Japanese Yen exchange rate and speculators' net positions



Notes: The grey line depicts the exchange rate of the Japanese Yen against the US Dollar, and the blue line the speculators' net positions.

Figure 1 plots the Yen and the large speculators' net positions in Yen futures at a weekly frequency. As we can see, the speculators' net positions tend to increase when the Yen is on an appreciation path, and they tend to decrease when the Yen is on a depreciation path. These patterns tend to be persistent: they may last for several months.

Figure 2. 9m-ahead predicted states and their predicted probabilities of Yen



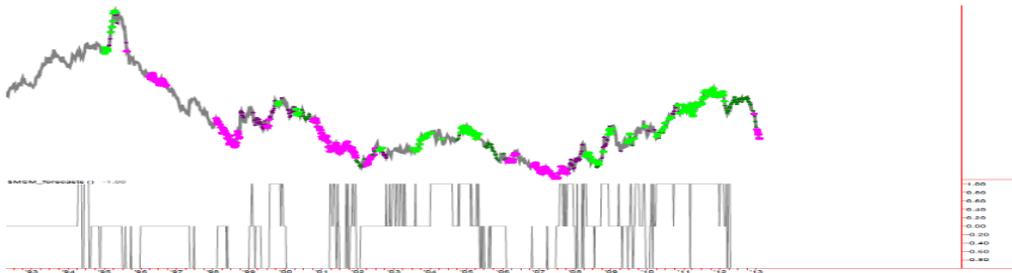
Notes: The predicted state takes values -1, 0 or 1. The value 1 (or -1) on a given week indicates that 9 months later, the most likely speculators's mode is accumulation (or decumulation). 0 means that the most likley mode is stasis.

Every week, we fit a three-state Markov switching model to the speculators' net positions. Every week, based on our MSM estimates, we form h -months

forecasts of the state. The forecast state \hat{S}_{t+h}^Y can take three values: -1 , 0 or 1 , which represent depreciation, no change and appreciation respectively. For example, Figure 2 shows the evolution of the 9-months ahead predicted state for the Yen (\hat{S}_{t+h}^Y). As we can see, the predicted state is quite persistent.

We then use the sequence of predicted states to construct our forecasts. Figure 3 illustrates the performance of these 9-month (i.e., 38 weeks) ahead directional forecasts with circles of different sizes and colors. If an appreciation forecast turned out to correctly predict the direction of the Yen between week $t-38$ and week t , then we place a big bright-green circle on week t 's exchange rate. In contrast, if the appreciation forecast was wrong, week t 's circle is small and dark-green. Similarly, if the depreciation forecast made 38 weeks ago turned out to correctly predict the subsequent Yen depreciation, then in week t there is a large bright-pink circle. Meanwhile, wrong depreciation forecasts are represented with a small dark-red circle.⁵

Figure 3. The performance of our 9m-ahead directional forecasts



Notes: The grey line depicts our 9m-ahead directional forecasts, which can be -1 , 0 or 1 . 1 (or -1) on a given week predicts appreciation (or depreciation) of the exchange rate Yen/US dollar. 0 means that our directional forecast is no change.

As we can see, 360 out of 571 circles are either bright-green or bright-pink.

⁵Readers looking at a black and white version of this paper can only see succesful forecasts (represented with large circles) and failed forecasts (represented with small circles).

This generates the 63 percent success ratio shown in Table 1. Interestingly, the successful directional forecasts in Figure 3 tend to be followed by bigger moves in the Yen than the moves following wrong directional forecasts. The tests we consider below evaluate formally whether such patterns could be generated by a random walk.

2 Literature Review

Our paper is linked to several branches of the exchange rate forecasting literature. First, our method of extracting information from the speculators' position data relies on the autoregressive Markov switching model proposed in Hamilton (1989, 1990). Engel and Hamilton (1990) apply this model to exchange rate data to explain the long swings exhibited by exchange rates from the mid 1970s to the end of the 1980s. Engel (1994), however, finds that this model does not clearly outperform the random walk in out-of-sample exchange rate forecasts. In this paper, the Markov switching model is applied to the speculators' positions, instead of the exchange rate itself. Furthermore, rather than a two-state MSM, we consider three regimes: uptrend, downtrend, and random walk.

By investigating the information content of market participants' trading positions, our paper is linked to the pioneering paper by Evans and Lyons (2002), which shows that the order flow of a group of dealers in the interbank foreign exchange market can forecast exchange rates over short horizons. While the order flow data they consider is private information, the COT data we use is public information.

According to recent survey papers by Cheung et al (2005), Rogoff and Stavrakeva (2008) and Rossi (2013), the Meese and Rogoff (1983) puzzle is still alive: at horizons of less than one year, the driftless random walk beats fundamentals-based exchange rate out-of-sample forecasts, with few exceptions. Using a panel specification, Engel, Mark and West (2007) test whether fundamentals-based models beat the random-walk. Using the CW-test, they find predictability over a 4-year horizon for the monetary model and the PPP

model in 11 and 13 out of 18 currencies, respectively. However, over 1-quarter they don't find predictability.⁶ Molodtsova and Papell (2009) consider the 1-month ahead predictability of Taylor rule based models for individual currencies. In their most successful specification—heterogenous coefficients, smoothing, and a constant—the CW-test rejects the null (at the 0.1 significance level) in favor of a symmetric Taylor rule model for 10 out of 12 currencies. Gourinchas and Rey (2007) use net foreign assets as a predictor of future exchange rate changes and forecast the trade-weighted US dollar rather than individual exchange rates. Using the CW-test they find that net foreign assets can predict the exchange rate better than the driftless random walk at both long and short horizons.⁷ Although, none of these papers carry out the DMW-test, Rossi (2013) reports that at both the 1-quarter and the 4-year forecasting horizons, the DMW-test finds no evidence of forecastability for the monetary model, the PPP model, and the Taylor rule based model across all currencies she considers.⁸ In this paper, we carry out both the CW and DMW tests. As we described in the Introduction, at the 6-to-12 months horizons, both tests reject the respective nulls in favour of our model in basically all currency-horizon pairs, except the Swiss Franc. However, at the 1 and 3 months forecasting horizons there is weak evidence of forecastability of our model.

There are many other papers that test the evidences of traditional predictors of exchange rate changes: interest rate differentials (UIRP) are considered by Meese and Rogoff (1988), Cheung, Chinn and Pascual (2005), Alquist and Chinn (2008) among others; price and inflation differentials (PPP) is considered in Rogoff (1996), Cheung, Chinn and Pascual (2005); Meese and Rogoff (1983), Chinn and Meese (1995), Mark(1995), Kilian (1999), Berkowitz and Giorgianni (2001), Faust et al (2003) and Rossi (2005) explore exchange rate forecastability using the monetary model. Alquist and Chinn (2008) use the net foreign asset model. Chen and Rogoff (2003), and Chen, Rogoff and Rossi

⁶Similar results are reported by Mark and Sul (2001).

⁷Net foreign assets is the deviation from trend of a weighted combination of gross assets, gross liabilities, gross exports and gross imports.

⁸She does not carry out the DMW test for the net foreign assets model.

(2010) exploit commodity price indices.

Like this paper, there are several papers that use commitment-of-traders data. Moskowitz et al (2012) document momentum over the time dimension across different asset classes and find that speculators profit from this time series momentum. Hong and Yogo (2012) find that open-interest in futures markets has predictive power over the excess returns of several assets classes over an one-month horizon. Brunnermeier and Petersen (2008) find that, in-sample, currency crash risk increases with speculators' positions in currency futures markets. Unlike our paper, out-of-sample forecasting is not the focus of these papers. Finally, Tornell and Yuan (2012) find that the peaks and troughs of speculators net positions are generally useful predictors of future exchange rates.

3 The Portfolio Problem and Its Empirical Implications

Our empirical strategy consists of extracting the implicit exchange rate forecasts from the speculators' net positions data in the COT report. To do so we derive an empirical Markov switching model from the portfolio problem of a representative investor in the foreign exchange futures market. In the next Section, we will bring this model to the data.⁹

3.1 The Portfolio Problem

Let e_t be the time- t price of a foreign currency futures contract, in terms of US Dollars. To simplify the exposition, we will refer to any non-US dollar

⁹The COT report provides weekly positions for three groups of traders in the US futures markets: hedgers, large speculators and the residual. Hedgers use the futures markets to insure against exchange rate changes, while large speculators participate in the futures markets to make capital gains. Even though the futures market is tiny compared with the spot foreign exchange market, speculators' positions contain valuable information about exchange rates. According to the 2013 BIS Triennial Central Bank Survey of turnover in foreign exchange markets, the daily average trading volume of spot markets worldwide is 2 trillion dollars in April 2013. Meanwhile, the value of foreign currency futures and options is around 160 billion dollars during the same period.

currency as a Euro. Thus, we will refer to e_t simply as the exchange rate and to an increase(decrease) in e_t as an appreciation(depreciation) of the Euro.

As in Engel and Hamilton (1990), the exchange rate is assumed to be driven by an unobservable stochastic trend x_t and a transitory component ε_t :

$$\Delta e_{t+1} = e_{t+1} - e_t = x_t + \varepsilon_{t+1}, \varepsilon_{t+1} \sim i.i.d. N(0, \sigma_\varepsilon^2) \quad (3.1)$$

where $N(0, \sigma_\varepsilon^2)$ denotes a normal random variable with mean zero and variance σ_ε^2 . The unobservable trend x_t is the sum of a state-dependent mean $\mu_{(S_t)}$ and an i.i.d. noise $u_t \sim N(0, \sigma_u^2)$, i.e.

$$x_t = \mu_{(S_t)} + u_t. \quad (3.2)$$

Depending on the state S_t , the mean $\mu_{(S_t)}$ of the stochastic trend x_t may take a positive, negative or zero value. The state S_t is a discrete first-order Markov-switching random variable with transition probability matrix $\Pi = [p_{i,j}]_{1 \leq i,j \leq 3}$, i.e.

$$\mu_{(S_t)} = \begin{cases} \mu_{(1)} > 0, & \text{if } S_t = 1 \text{ (up-trend)} \\ \mu_{(2)} = 0, & \text{if } S_t = 2 \text{ (range)} \\ \mu_{(3)} < 0, & \text{if } S_t = 3 \text{ (down-trend)} \end{cases}, \quad (3.3)$$

where $p_{i,j} = \Pr(S_t = j | S_{t-1} = i)$ is the conditional probability that the state is j at time t , given that it was i at time $t - 1$. We take the unobservable trend x_t as a primitive. It might depend on expectations of fundamentals, sentiment or other unobservable factors.

There are overlapping generations of two-period lived risk-averse speculators. A young t -date speculator observes a noisy signal (y_t) of the unobservable exchange rate trend x_t :

$$y_t = x_t + v_t, v_t \sim i.i.d. N(0, \sigma_v^2). \quad (3.4)$$

This signal, allows the speculator to form better forecasts of future exchange rate changes than those based on the historical exchange rate data she ob-

serves. This is because the unobservable trend x_t affects Δe_{t+1} directly and the information of x_t is directly contained in y_t . Meanwhile, with available exchange rate data, i.e. $\{\Delta e_s\}_{s=1}^t$, the speculator can only estimate x_{t-1} .¹⁰

During the first period of her life, the young representative speculator takes a position b_t in the Euro futures market. During the second period of her life, the now old speculator closes her position. The representative young speculator is risk averse and solves the following problem

$$\max_{b_t \in \mathbb{R}} E[-\exp(-\gamma W_{t+1}) | I_t], \text{ with } W_{t+1} = b_t(e_{t+1} - e_t) \quad (3.5)$$

where $I_t = \{y_1, \dots, y_t\}$ denotes the information at time t , γ is the risk-aversion coefficient and W_{t+1} is her wealth next period.¹¹ The speculator may take either a long position ($b_t > 0$), a short position ($b_t < 0$) or stay out of the market ($b_t = 0$). Notice that our model's speculator faces no position limits and she needs not post margin.¹²

The representative date- t young speculator's prior distribution on the stochastic exchange rate trend x_t is given by the posterior distribution of her predecessor. To close the model, we assume that the prior of the first cohort (date-1) of young speculators is $x_0 \sim N(0, \sigma_u^2)$.

In order to solve the t -date young representative speculator's problem, note that she enters period t with a prior $\hat{x}_{t|t-1} \sim N(\hat{x}_{t-1}, \sigma_{t-1}^2)$, where \hat{x}_{t-1} and σ_{t-1}^2 are the mean and variance of her predecessor's posterior distribution. Bayesian updating implies that after observing the signal y_t , the t -date speculator's posterior on x_t is normal with mean \hat{x}_t and variance σ_t^2 , which are given by

¹⁰The signal y_t can be thought of as the speculator's forecasts of the exchange rates, which is unobservable to the econometrician.

¹¹We treat b_t as a real number, while the futures contract in the Chicago Mercantile Exchange is for 125,000 Euros.

¹²To a first-order this is a realistic assumption as the margin requirements are quite small. For instance, on February 7, 2013 the initial margin required on a Euro futures contract at the CME was around \$2500, while the dollar value of the Euro futures contract was \$167,000. Notice also that the payoff does not include the interest rate differential because it is embedded in the futures' price spread Δe_{t+1} .

the following filter

$$\hat{x}_t = (1 - k_t)\hat{x}_{t-1} + k_t y_t \quad (3.6)$$

$$\sigma_t^2 = k_{t-1}\sigma_v^2 \text{ and } k_{t-1} = \frac{\sigma_{t-1}^2 + \sigma_u^2}{\sigma_{t-1}^2 + \sigma_u^2 + \sigma_v^2}. \quad (3.7)$$

Since the prior of the first cohort (date-1) of speculators is $x_0 \sim N(0, \sigma_u^2)$, this recursion is initialized at $\hat{x}_0 = 0$ and $\sigma_0^2 = \sigma_u^2$.

It follows from (3.1) and (3.6) that the speculator's posterior belief is that exchange rate changes are normal with mean $E_t[\Delta e_{t+1}] = \hat{x}_t$ and variance $\text{Var}_t(\Delta e_{t+1}) = \sigma_t^2 + \sigma_\varepsilon^2$. Thus, the speculator's problem (3.5) can be rewritten as

$$\max_{b_t \in \mathbb{R}} \left[-\exp \left(-\gamma \hat{x}_t b_t + \frac{\gamma^2}{2} (\sigma_t^2 + \sigma_\varepsilon^2) b_t^2 \right) \right].$$

By taking the derivative with respect to b_t , we have that the representative speculator's demand for Euro futures is

$$b_t^* = \frac{E(\Delta e_{t+1} | I_t)}{\gamma \cdot \text{Var}(\Delta e_{t+1} | I_t)} = \frac{\hat{x}_t}{\gamma(\sigma_t^2 + \sigma_\varepsilon^2)}. \quad (3.8)$$

Equation (3.8) says that the representative speculator's Euro position is increasing in her posterior about the Euro trend (\hat{x}_t) and it is decreasing in the degree of risk-aversion and the variance of expected returns.

3.2 The Empirical Model

The empirical counterpart of the model's demand for Euro futures b_t^* is the speculators' net position data in the COT report, which we denote by Z_t . That is, $Z_t = b_t^*$.

Here, we use the solution to the speculator's problem (equations (3.6)-(3.8)) to recover the information about y_t —the noisy signal observed by speculators—from the observed data on the speculators' net positions Z_t . These equations show that if speculators observe a long enough sequence of positive(negative) y_t -signals, they gradually increase their belief that the Euro will appreci-

ate(depreciate). Based on these posterior beliefs speculators gradually increase their long(short) Euro position b_t^* .

Note from the first equation in the speculators' filter (3.6) that we can express the y_t -signal as follows

$$y_t = \frac{1}{k_t}\hat{x}_t - \frac{1-k_t}{k_t}\hat{x}_{t-1} = g_t\hat{x}_t + (1-g_t)\hat{x}_{t-1}, \text{ where } g_t \equiv k_t^{-1}. \quad (3.9)$$

Using the representative speculator's demand for Euro futures (3.8) $\hat{x}_t = \gamma(\sigma_t^2 + \sigma_\varepsilon^2)b_t^*$. Replacing this expression in (3.9), we see that the unobserved signal can be expressed in terms of speculators' net positions as follows

$$y_t = \gamma(\sigma_t^2 + \sigma_\varepsilon^2)g_t b_t^* + \gamma(\sigma_{t-1}^2 + \sigma_\varepsilon^2)(1-g_t)b_{t-1}^*. \quad (3.10)$$

Since $Z_t = b_t^*$, equation (3.10) implies that the speculators' net positions follow an AR(1) process

$$Z_t = \theta_t Z_{t-1} + \sigma_{x,t} y_t, \quad (3.11)$$

where

$$\sigma_{x,t} = \frac{1}{\gamma(\sigma_t^2 + \sigma_\varepsilon^2)g_t} \text{ and } \theta_t = \frac{\sigma_{t-1}^2 + \sigma_\varepsilon^2}{\sigma_t^2 + \sigma_\varepsilon^2}(1-k_t). \quad (3.12)$$

In the appendix, we show that σ_t^2 is convergent, such that the limit of θ_t is in $(0, 1)$.

4 Estimation of the Empirical Model

We use model (3.11) to recover the signal observed by speculators by estimating, for each currency, the following first order autoregressive (AR(1)) MSM using rolling samples:

$$Z_t = \theta Z_{t-1} + \mu_{(S_t)}^* + v_t^*, \text{ with } \mu_{(S_t)}^* = \sigma_x \mu_{(S_t)} \text{ and } v_t^* = \sigma_x(v_t + u_t), \quad (4.1)$$

where the Markov switching component $\mu_{(S_t)}$ is defined in (3.3) and v_t^* is i.i.d. normal with mean zero and variance $\sigma_{v^*}^2$.¹³ We allow for heterogeneous DGPs for the speculators' positions $Z_{i,t}$ in different currencies, indexed by i . For ease of notation, however, we ignore the subscript 'i' in $Z_{i,t}$, unless necessary.

Given the unknown parameters $\alpha = (\theta, \mu_{(1)}^*, \mu_{(2)}^*, \mu_{(3)}^*, \sigma_{v^*}^2, \Pi)$ and observations on the speculators' net positions up to time t : $\mathbf{Z}^t = (Z_t, Z_{t-1}, \dots, Z_0)$, the density of Z_t conditional on the state S_t taking the value j is given by

$$f(Z_t | \mathbf{Z}^{t-1}, S_t = j; \alpha) = \frac{1}{\sqrt{2\pi}\sigma_{v^*}} \exp \left[-\frac{(Z_t - \theta Z_{t-1} - \mu_{(j)}^*)^2}{2\sigma_{v^*}^2} \right], \quad j = 1, 2, 3. \quad (4.2)$$

Given the prediction probabilities

$$\xi_{j,t-1} = \Pr(S_t = j | \mathbf{Z}^{t-1}; \alpha), \quad j = 1, 2, 3 \quad (4.3)$$

we can calculate the conditional density of Z_t given $\mathbf{Z}^{(t-1)}$ as

$$f(Z_t | \mathbf{Z}^{(t-1)}; \alpha) = \sum_{j=1}^3 f(Z_t | S_t = j, \mathbf{Z}^{t-1}; \alpha) \xi_{j,t-1}. \quad (4.4)$$

Using the above conditional density we can compute filtered probabilities

$$\Pr(S_t = j | \mathbf{Z}^t; \alpha) = \frac{f(Z_t | \mathbf{Z}^{t-1}, S_t = j; \alpha) \xi_{j,t-1}}{f(Z_t | \mathbf{Z}^{t-1}; \alpha)} \quad \text{for } j = 1, 2, 3, \quad (4.5)$$

which together with the transition probabilities implies that the filtered probabilities are

$$\xi_{j,t} = \sum_{k=1}^3 p_{k,j} \Pr(S_t = k | \mathbf{Z}^t, \alpha) \quad \text{for } j = 1, 2, 3. \quad (4.6)$$

¹³The theoretical model (3.11)-(3.12) indicates that Z_t follows an AR(1) process with time varying coefficients $(\theta_t, \sigma_{x,t})$. The rolling sample estimation of model (4.1) can be viewed as a local constant approximation of the theoretical model (3.11)-(3.12). In Appendix A, we show that $(\theta_t, \sigma_{x,t})$ has a convergent fixed point (θ^*, σ_x^*) with $|\theta^*| < 1$.

The log-likelihood function is therefore

$$Q_n(\alpha) = \sum_{t=1}^n \log [f(Z_t | \mathbf{Z}^{t-1}, \alpha)],$$

where, given α , $f(Z_t | \mathbf{Z}^{t-1}, \alpha)$ is calculated using (4.2)-(4.6) with the initial values $\Pr(S_1 = k | \mathbf{Z}^0, \alpha)$ for $t = 1, \dots, n$.

The log-likelihood function $Q_n(\alpha)$ is a highly nonlinear function of α . Thus, we use the EM algorithm proposed by Hamilton (1990) to obtain the maximum likelihood (ML) estimator $\hat{\alpha}_n$ of α .¹⁴ To ensure that a global maximum is attained, we consider 150 different initial values for maximum likelihood estimation. Given the ML estimator $\hat{\alpha}_n$, we estimate the filtered probabilities by $\Pr(S_t = j | \mathbf{Z}^t, \hat{\alpha}_n)$ and the prediction probabilities by $\Pr(S_{t+1} = j | \mathbf{Z}^t, \hat{\alpha}_n)$ for $t = 1, \dots, n$.

4.1 Data

We obtain data on the large speculators' positions and open interest from the COT Report published by the CFTC. Typically, every week the CFTC gathers positions data on Tuesday and publishes the COT report three days later, on Friday after markets close. Thus, for each of the six currencies we consider, every Friday we fit Markov switching model (4.1) to the speculators' net positions normalized by open interest. Based on these estimates we generate our new forecasts every Friday. We evaluate our forecasts using weekly spot exchange rates at the end of trading on Friday, released by the Federal Reserve Board.

Our sample begins in September 1992, which is when the CFTC started to release the COT report on a weekly basis, except for the Euro, for which the data begins in January 1999. There are a few missing values in net positions and open interest data over the sample period.¹⁵ We replace the missing values

¹⁴The main advantage of the EM algorithm over direct numerical optimization methods is its robustness with respect to the multiple local maxima problem.

¹⁵There are 27 missing values for the Australian Dollar and 2 for the Swiss Franc.

using a linear interpolation, using the last available value before the missing value and the first available value after the missing value.

For each of the six currencies ($i = 1, \dots, 6$), we estimate the Markov switching model (4.1) using an initial sample with $n_0^{(i)}$ observations, i.e. $\{Z_{i,t}\}_{t=1}^{n_0^{(i)}}$. We then re-estimate the Markov switching model using a rolling sample: Every week we add a new observation and drop the first observation in the previous sample. Thus, if for currency i we have $n^{(i)}$ observations, the model is estimated $n^{(i)} - n_0^{(i)}$ many times. We consider three rolling window sizes $\{80, 100, 120\}$ weeks. For each currency, we choose the rolling window size that generates the highest average forecast success ratio across the five forecasting horizons we consider: $1m$, $3m$, $6m$, $9m$, and $12m$.¹⁶ We thus set the rolling window size and the initial sample to 80 for the Japanese Yen and Swiss Franc; 100 for the Canadian Dollar; and 120 for the Euro, British Pound and Australian Dollar.

Figure D.1 in the appendix depicts the series of the three estimated means $\{\mu_{(1)}^*, \mu_{(2)}^*, \mu_{(3)}^*\}$ of each state in our AR(1) Markov switching model (4.1). As we can see, for all currencies the means of the up-trend states $\mu_{(1)}^*$ are strictly positive; the means of the down-trend states $\mu_{(3)}^*$ are negative; while the means of the range states $\mu_{(2)}^*$ fluctuate around zero. Figure D.2 in the appendix depicts the series of the estimated autoregressive coefficient θ for each currency. Its value ranges between 0 and 1, which is consistent with our theoretical model (see equation (3.11)).

5 Directional Forecasts of the Exchange Rate

We construct our exchange rate forecasts in two steps. In the first step, we use the estimates of model (4.1) to compute the most likely accumulation mode of the speculators in each period between t_0 and $t_0 + h$. In the second step, we use this sequence of predicted states to obtain our exchange rate forecasts. Specifically, we forecast appreciation between t_0 and $t_0 + h$ if in a majority

¹⁶The forecast success ratios are considered in the next subsection.

of periods, accumulation is the most likely speculators' mode. Conversely, we forecast depreciation if decumulation is the most likely speculators' mode. In the other cases, we stick to the random walk forecasts and predict no exchange rate change over an h -horizon.

In the first step, we compute the h -week ahead prediction probabilities of each state for each currency i , based on t_0 information, as follows

$$\left(\widehat{\xi}_{1,t_0+h}^{(i)}, \widehat{\xi}_{2,t_0+h}^{(i)}, \widehat{\xi}_{3,t_0+h}^{(i)} \right) = \left(\widehat{\xi}_{1,t_0}^{(i)}, \widehat{\xi}_{2,t_0}^{(i)}, \widehat{\xi}_{3,t_0}^{(i)} \right) (\widehat{\Pi}_{t_0}^{(i)})^h, \quad (5.1)$$

where $\widehat{\Pi}_{t_0}^{(i)}$ is the estimated transition matrix based on the rolling sample of speculators' net positions $\mathbf{Z}_i^{(t_0)} = (Z_{i,t_0-n_0+1}, \dots, Z_{i,t_0})$, and $\widehat{\xi}_{j,t_0}^{(i)}$ is the estimated filtered probabilities state j . It is given by

$$\widehat{\xi}_{j,t_0}^{(i)} = P(S_{t_0}^{(i)} = j | \mathbf{Z}_i^{(t_0)}, \widehat{\alpha}_{n_0}^{(i)}(t_0)) \text{ for } j = 1, 2, 3, \quad (5.2)$$

where $\widehat{\alpha}_{n_0}^{(i)}(t_0)$ is the estimator of the unknown parameters based on the rolling sample $\mathbf{Z}_i^{(t_0)}$.

Based on the prediction probabilities (5.1), our time- t_0 prediction of the most likely speculators' mode in period $t_0 + h$, which we will refer to as the "predicted state," is given by

$$\widehat{S}_{t_0+h}^{(i)} = \begin{cases} 1, & \text{if } \widehat{\xi}_{1,t_0+h}^{(i)} > \max \left\{ \widehat{\xi}_{2,t_0+h}^{(i)}, \widehat{\xi}_{3,t_0+h}^{(i)} \right\} \\ 2, & \text{if } \widehat{\xi}_{2,t_0+h}^{(i)} > \max \left\{ \widehat{\xi}_{1,t_0+h}^{(i)}, \widehat{\xi}_{3,t_0+h}^{(i)} \right\} \\ 3, & \text{if } \widehat{\xi}_{3,t_0+h}^{(i)} > \max \left\{ \widehat{\xi}_{1,t_0+h}^{(i)}, \widehat{\xi}_{2,t_0+h}^{(i)} \right\} \end{cases}. \quad (5.3)$$

$\widehat{S}_{t_0+h}^{(i)} = 1$ (resp. 3) means that the time- t_0 forecast of the most likely speculators' mode on currency i at time $t_0 + h$ is accumulation (resp. decumulation).

In the second step, the sequence of predicted states $\left\{ \widehat{S}_{t_0+1}^{(i)}, \widehat{S}_{t_0+2}^{(i)}, \dots, \widehat{S}_{t_0+h}^{(i)} \right\}$ are used to construct our exchange rate directional forecasts. Specifically, at t_0 we forecast an exchange rate appreciation (depreciation) over the following h periods, if between t_0 and $t_0 + h$ the majority of the predicted states indicate speculators' accumulation (decumulation). In other cases, we forecast zero ex-

change rate change. That is, for any t_0 with $n_0^i \leq t_0 \leq n - h$, our directional exchange rate forecasts are:

$$D_{t_0,h}^{(i)} = \begin{cases} 1, & \text{if } X_{t_0,h}^{(i)} > 0 \\ 0, & \text{if } X_{t_0,h}^{(i)} = 0 \\ -1, & \text{if } X_{t_0,h}^{(i)} < 0 \end{cases}, \quad X_{t_0,h}^{(i)} \equiv \sum_{s=1}^h \left(2 - \widehat{S}_{t_0+s}^{(i)} \right) \quad (5.4)$$

Notice that $D_{t_0,h}^{(i)} = 1$ (resp. -1) means that our directional forecast over an h -period horizon is appreciation (resp. depreciation). Meanwhile, if $D_{t_0,h}^{(i)} = 0$, we predict no change. The variable $X_{t,h}^{(i)}$ can take values ranging from $-h$ to $+h$. It indicates the net number of periods with predicted speculators' accumulation between t_0 and $t_0 + h$ (if $X_{t,h}^{(i)} > 0$) or the net number of periods with predicted decumulation (if $X_{t,h}^{(i)} < 0$). For example, $X_{t,h}^{(i)} = -\omega$ when speculators' decumulation is the predicted state in ω periods on $(t, t+h)$; while $X_{t,h}^{(i)} = +\omega$ when speculators' accumulation is the predicted state in ω periods on $(t, t+h)$.

Intuitively, one can think of the transition probabilities as capturing the low frequency component of the speculators' accumulation mode, whereas the filtered probabilities capture the high frequency component of the speculators mode. Equation (5.1) implies that for a short forecasting horizon h , the predicted probabilities of the state $\widehat{\xi}_{j,t_0+h}^{(i)}$ are determined mainly by the time- t filtered probabilities $\widehat{\xi}_{j,t_0}^{(i)}$. In contrast, for long enough h the predictions are determined by the ergodic distribution implied by the transition matrix $\widehat{\Pi}_{t_0}^{(i)}$, by the well known convergence property of Markov chains. Our directional exchange rate forecasts in (5.4) have the same properties: at short horizons they are determined mainly by the estimated filtered probabilities, whereas at long horizons they are determined mainly by the estimated transition matrix.

During every week, we generate out-of-sample exchange rate directional forecasts for five horizons: $1m$, $3m$, $6m$, $9m$, and $12m$. For each currency i , our directional forecasts start the week after the first estimation of the MSM, and end h weeks prior to the end of our sample (week n). That is, they start in week $n_0^{(i)}$ and end in week $n - h$. Our sample starts on 10/2/1992

for all currencies except the Euro, which starts on 01/08/1999. Our sample ends on 02/08/2013 for all currencies. Thus, our directional forecasts start on 01/20/1995 for the Australian Dollar and British Pound; 09/01/1994 for the Canadian Dollar; 04/15/1994 for the Swiss Franc and Japanese Yen; and 05/04/2001 for the Euro.¹⁷

Table 1 exhibits the forecast success ratio of our out-of-sample directional forecasts at the $1m$, $3m$, $6m$, $9m$, and $12m$ horizons.¹⁸ The forecast success ratio is the number of successful depreciation or appreciation forecasts divided by the total number of depreciation and appreciation forecasts. As we can see, the aggregate forecast success ratio is 58 percent over the period October 1992-February 2013. Interestingly, our forecast accuracy is greater at the $6m$ to $12m$ horizons than at the $1m$ to $3m$ months horizons. If we confine our attention to $6m$, $9m$ and $12m$ forecasting horizons, the aggregate success ratio is 60 percent (64 percent excluding the Swiss Franc). When taken individually, we can see that the $6m$ to $12m$ success ratios are greater than 56 percent in all cases—except for the Swiss Franc—and that in several cases the success ratio is larger than 66 percent.

The high average success ratios observed in Table 1 are not dominated by specific periods of successful forecasts. In most country-horizon pairs, the performance of the directional forecast is stable over the sample period. Figures D.3 through D.7 in the appendix exhibit this stability by plotting the evolution of the cumulative forecast success ratios for each forecasting horizon. As we can see, approximately after 400 weeks, these ratios converge to a stable level, above 50 percent in most country-horizon pairs. Initially, however, these ratios fluctuate quite a bit because the sample size is small. For example, in Figure D.6 the $9m$ ahead forecasts for the yen initially fluctuate between 60 percent and 90 percent before converging to the mid-60 percents.

Table 2 considers the same forecast success ratios as those considered in

¹⁷The dates of the last forecasts are 01/18/2013 for $h = 1m$, 11/16/2012 for $h = 3m$, 08/24/2012 for $h = 6m$, 05/25/2012 for $h = 9m$, and 03/01/2012 for $h = 12m$.

¹⁸Throughout the paper 1, 3, 6, 9 and 12 months correspond to 4, 13, 25, 38 and 50 weeks, respectively.

Table 1, but partitions the sample period into two sub-periods: 04/15/1994-11/29/2002 and 12/06/2002-02/08/2013. As we can see, the forecast success ratios are very similar across the two sub-periods for the Yen, the British Pound, Australian Dollar and the Swiss Franc. For the Canadian Dollar, however, the ratios are around 20 percent higher in the second period.

5.1 Evaluating the Statistical Significance of Directional Forecasts

The high forecasting success ratios in Table 1 illustrate the usefulness of the COT data in generating directional exchange rate forecasts. The issue remains, however, whether this success isn't simply the result of luck. Wouldn't flipping a coin result in similar success ratios? Here, we investigate the statistical significance of the success of our directional forecasts by conducting two formal tests of the null hypothesis that our MSM-based forecasts cannot improve the directional forecasts relative to the driftless random walk forecasts.

First, we propose a novel test that weights each directional forecast by the realized exchange rate change, and so gives more weight to the directional forecasts associated bigger exchange rate moves. Second, we consider the standard binomial test, which gives the same weight to all forecasts.¹⁹

The forecasting evaluation criterion of our weighted test is nested by the generic lose function studied in Diebold and Mariano (1996), West (1996) and others. However, its special form makes it more relevant to situations faced by real world investors as it is linked to the profitability of trading strategies. It captures the spirit of George Soros's observation: It's not whether you're right or wrong, but how much money you make when you're right and how much you lose when you're wrong. The statistic in our weighted directional test has the same spirit as the trading rule considered by Elliott and Ito (1999) to evaluate the predictability content of survey forecasts. They use survey forecasts of the yen/dollar exchange rate in their rule, and find that the resulting profit is positive, but small.

¹⁹For instance, Pesaran and Timmerman (1992), Leitch and Tanner (1991), Engel (1994) and Cheung, et.al. (2005). See Elliot and Timmermann (2008) for a discussion of directional tests.

Figures D.8 through D.12 in the appendix depict the evolution of our directional forecasts $D_{t,h}^{(i)}$ in (5.4) together with the actual exchange rate movement for all currencies.

5.1.1 Weighted Directional Test

Here, we test the null hypothesis that our weighted directional forecasts are unable to improve upon the driftless random walk model:

$$e_{t+1}^{(i)} = e_t^{(i)} + \varepsilon_{t+1}^{(i)}, \quad (5.5)$$

where $\{\varepsilon_t^{(i)}\}$ is a white noise process with mean zero and variance $\sigma_{i,\varepsilon}^2$. We maintain the assumption that, under the null, $\{\varepsilon_t\}$ is a martingale difference sequence. Given the observations $\{e_{t+h}^{(i)} - e_t^{(i)}\}_{t=n_0^{(i)}}^{n^{(i)}-h}$ and our directional forecasts $\{D_{t,h}^{(i)}\}_{t=n_0^{(i)}}^{n^{(i)}-h}$ on each currency i , we consider the following test statistic

$$T_{a,n}^{(i)} = \frac{1}{n_1^{(i)}} \sum_{t=n_0^{(i)}}^{n^{(i)}-h} D_{t,h}^{(i)} (e_{t+h}^{(i)} - e_t^{(i)}), \text{ where } n_1^{(i)} \equiv n^{(i)} - n_0^{(i)} - h + 1. \quad (5.6)$$

Under the driftless random walk specification, the optimal forecast is a zero exchange rate change. If we replace $D_{t,h}^{(i)}$ in $T_{a,n}^{(i)}$ by the optimal forecasts generated under the driftless random walk hypothesis, $T_{a,n}^{(i)}$ becomes zero for all n . Thus, the null hypothesis underlying the test statistic $T_{a,n}^{(i)}$ can be specified as

$$H_0 : \quad E \left[D_{t,h}^{(i)} (e_{t+h}^{(i)} - e_t^{(i)}) \right] = 0 \text{ for any } i \text{ and } t. \quad (5.7)$$

That is, under the null our directional forecasts are uncorrelated with future realized exchange rate changes.

Notice that the above zero expectation holds as $\{\varepsilon_t^{(i)}\}$ is a martingale difference sequence with respect to the natural sigma field generated by the COT process $\{Z_{i,t}\}_t$ and the variables $D_{t,h}^{(i)}$ only depend on the COT data up to time t . On the other hand, if the COT data was useful in forecasting $e_{t+h}^{(i)} - e_t^{(i)}$, then one would expect that $T_{a,n}^{(i)}$ converge to some strictly positive number.

In order to test null hypothesis (5.7), let $V_{T_a^{(i)},n}$ denote the consistent estimator of the asymptotic variance of $T_{a,n}^{(i)}$. Then by Slutsky's theorem and the martingale central limit theorem, we deduce that

$$\sqrt{n_1^{(i)}} V_{T_a^{(i)},n}^{-1/2} T_{a,n}^{(i)} \rightarrow_d N(0, 1). \quad (5.8)$$

For the empirical implementation, we construct $V_{T_a^{(i)},n}$ using two LRV estimators that control for auto-correlation in $\left\{ D_{t,h}^{(i)} (e_{t+h}^{(i)} - e_t^{(i)}) \right\}_{t=n_0}^{n-h}$: the Newey-West LRV estimator and the Andrews LRV estimator.

Table 3 present the values of the $T_{a,n}^{(i)}$ statistic, and its t-values (test statistics), for the 6 currencies (Euro, Japanese Yen, British Pound, Australian Dollar, Canadian Dollar and Swiss Franc) and the 5 horizons ($1m$, $3m$, $6m$, $9m$ and $12m$) we consider. The null (5.7) tested is that our directional forecasts are uncorrelated with future realized exchange rate changes. The null is rejected if the test statistic $T_{a,n}^{(i)}$ is significantly larger than zero. For the one-sided test, a t -value greater than 1.282 implies a 0.1 significance level.

The high forecast success ratios in Table 1 translate into strong predictability of future exchange rate changes (except for the Swiss Franc) using our directional forecasts. Even after controlling the autocorrelation of the data, over the 6m to 12m horizons, there is very strong evidence of exchange rate predictability in all currencies, except for the Swiss Franc. As we can see in panel A of Table 3, using the Newey-West LRV estimator, the null is rejected in all the 15 country-horizon pairs at the $6m$, $9m$ and $12m$ horizons. Panel B shows that using the Andrews LRV estimator to control for autocorrelation, the null is rejected in 12 out of those 15 country-horizon pairs. At the $1m$ and $3m$ horizons there is weak evidence of predictability: the random walk null is rejected in only 5 out of 12 currency-horizon pairs.

5.1.2 Binomial Directional Test

Here, we test the significance of our model in forecasting the sign of $e_{t+h}^{(i)} - e_t^{(i)}$. When the exchange rate is a driftless random walk, we define a new dummy

variable $R_{t,h}^{(i)}$ that captures the direction of the realized exchange rate change of currency i over horizon h :

$$R_{t,h}^{(i)} = \begin{cases} 1, & \text{if } e_{t+h}^{(i)} - e_t^{(i)} \geq 0 \\ -1, & \text{if } e_{t+h}^{(i)} - e_t^{(i)} < 0 \end{cases}. \quad (5.9)$$

The null hypothesis we test is that our directional forecasts $D_{t,h}^{(i)}$ are uncorrelated with the future direction of the exchange rate $R_{t,h}^{(i)}$

$$H_0: Cov\left(D_{t,h}^{(i)}, R_{t,h}^{(i)}\right) = 0, \quad (5.10)$$

while the alternative hypothesis could be one sided or two sided:

$$H_1^{\text{one-sided}}: Cov\left(D_{t,h}^{(i)}, R_{t,h}^{(i)}\right) > 0 \text{ or } H_1^{\text{two-sided}}: Cov\left(D_{t,h}^{(i)}, R_{t,h}^{(i)}\right) \neq 0.$$

Consider then the following test statistic

$$T_{b,n}^{(i)} = \frac{1}{n_1^{(i)}} \sum_{t=n_0^{(i)}}^{n-h} D_{t,h}^{(i)} R_{t,h}^{(i)} - \frac{1}{n_1^{(i)}} \sum_{t=n_0^{(i)}}^{n-h} D_{t,h} \frac{1}{n_1^{(i)}} \sum_{t=n_0^{(i)}}^{n-h} R_{t,h}^{(i)},$$

which is the sample covariance of the two random variables: $D_{t,h}^{(i)}$ and $R_{t,h}^{(i)}$. In order to test null hypothesis (5.10), let $V_{T_{b,n}^{(i)}}$ denote the consistent estimator of the asymptotic variance of $T_{b,n}^{(i)}$. Then we have

$$\sqrt{n_1^{(i)}} V_{T_{b,n}^{(i)}}^{-\frac{1}{2}} T_{b,n}^{(i)} \rightarrow_d N(0, 1).$$

Overall the test results are very similar to those of the weighted directional test. Over $6m$, $9m$ and $12m$ forecasting horizons, the binomial test suggests that our directional forecasts have strong predictability for the directional changes in future exchange rate in 5 currencies, except the Swiss Franc.

Table 4 reports the test results using Newey-West LRV estimator (Panel A) and Andrews LRV estimator (Panel B) to control for auto-correlation. The

null is rejected if the t -value of the $T_{b,n}^{(i)}$ statistic is positive and statistically significant. For the one-sided test, a t -value greater than 1.282 implies a 0.1 significance level.

Excluding the Swiss Franc, the null is rejected in 17 (Panel A) and 16 (Panel B) out of the 25 currency-horizon pairs. Over $6m$, $9m$ and $12m$ horizons, strong evidence of directional predictability is found in 14 (Panel A) and 13 (Panel B) out of 15 cases, excluding the Swiss Franc.

6 Evaluating the Accuracy of Point Forecasts

The standard practice in the exchange rate forecasting literature has been to test the accuracy of out-of-sample point-forecasts of various models vis-a-vis random-walk forecasts. The most widely used tests are those proposed by Diebold and Mariano (1995) and West (1996)—the DMW test—and by Clark and West (2006)—the CW test. In this section, we use our MSM to generate h -period-ahead point forecasts and carry out the standard DMW and CW tests against the driftless random walk, which has proven to be a tougher benchmark to beat than the random walk with drift. In Section C we consider the random walk with drift.

We are aware of the potential degeneracy issue when one model nests the other, which is the case in the comparison of our forecasts with those of the driftless random walk model. In this case, the asymptotic variance of the DMW test statistic is zero when the estimation error of the model’s parameters vanishes fast as the sample size grows. In practice, however, the estimation error of the model’s parameters could be nontrivial, because the samples used to estimate the model’s parameters and generate the forecasts of the state variables are relatively small. In such scenario, the asymptotic normal approximation to DMW test statistic may still provide valid inference in practice (see, e.g., Giacomini and White (2006)). On the other hand, our CW test takes the model degeneracy into account. It corrects the quadratic term of parameter estimation errors in the DMW test statistic, and conducts inferences based on the normal approximation of the errors in estimating the expectation of the

forecasting evaluation loss function. Our motivation for conducting both the DMW and the CW tests is to check the robustness of the significance of our empirical results.

Our h -period ahead point forecast for currency i is that between t and $t+h$, the magnitude of the appreciation(depreciation) will be proportional to the net number of periods with predicted speculators' accumulation(decumulation) over the following h -weeks, given by $X_{t,h}^{(i)}$ in (5.4). That is, the h -period ahead point forecasts for currency i , made at time t , is

$$\widehat{e}_{t+h}^{(i)} = e_t^{(i)} + \widehat{\beta}_{h,m_0} X_{t,h}^{(i)}, \text{ for } m_0 + 1 \leq t \leq n^{(i)} - h, \quad (6.1)$$

where m_0 is the first week that the out-of-sample point forecasts begin, $n^{(i)}$ is the last week of the speculator's net position data sample, and $\widehat{\beta}_{h,m_0}$ denotes the estimated effect of $X_{t,h}^{(i)}$ on the exchange rate change over horizon h . It is given by

$$\widehat{\beta}_{h,m_0} = \frac{\sum_{i \in I} \sum_{t=n_0^{(i)}+1}^{m_0-h} X_{t,h}^{(i)} (e_{t+h}^{(i)} - e_t^{(i)})}{\sum_{i \in I} \sum_{t=n_0^{(i)}+1}^{m_0-h} [X_{t,h}^{(i)}]^2}, \quad I = \{ad, bp, cd, jy, sf\}. \quad (6.2)$$

Three comments are in order. First, the summation in (6.2) starts in $n_0^{(i)} + 1$ because we use the initial $n_0^{(i)}$ weekly COT observations to estimate the MSM for currency i . Second, while for five currencies our sample period starts on October 2, 1992, for the Euro it starts on January 8, 1999. Thus, to maximize the number of out-of-sample forecasts, we exclude the Euro from the estimation of $\widehat{\beta}_{h,m_0}$. Third, we estimate $\widehat{\beta}_{h,m_0}$ using OLS, setting $m_0 = 360$. In this way, we synchronize to the 360th week the time when the first point forecast is generated across all currencies, except the Euro. It follows that our first out-of-sample point forecasts start on 04/09/1999 for the five currencies excluding the Euro. For the latter it starts on 05/04/2001. Our last point forecasts are generated on February 8, 2013. Figures D.13 through D.17 in the appendix depict the evolution of the net number of periods with predicted speculators' accumulation (decumulation) over the following h -weeks $X_{t,h}^{(i)}$ in

(5.4) together with the actual exchange rate movement for all currencies.

6.1 The Diebold-Mariano-West and the Clark-West Tests

The test proposed by Diebold and Mariano (1995) and West (1996) tests the null that the mean squared prediction error (MSPE) of a random walk is equal to the MSPE generated by the point forecasts of a given model. Meanwhile, the test proposed by Clark and West (2006) tests the null that a forecasting model is equivalent to the random walk model. In this section, we implement both tests.

Under the assumption that the exchange rate follows a driftless random walk, the h -period-ahead point forecast for currency i is $\bar{e}_{t+h}^{(i)} = e_t^{(i)}$ ($i = 1, \dots, 6$). Thus, to carry out the DWM test, we consider the following $DMW_{h,n}^{(i)}$ test statistic

$$DMW_{h,n}^{(i)} = \frac{1}{n_1^{(i)}} \sum_{t=m_0+1}^{n^{(i)}-h} \left[\left(e_t^{(i)} - e_{t+h}^{(i)} \right)^2 - \left(\hat{\beta}_{h,m_0} X_{t,h}^{(i)} + e_t^{(i)} - e_{t+h}^{(i)} \right)^2 \right], \quad (6.3)$$

where $n_1^{(i)} = n^{(i)} - m_0 - h$. That is, our point forecasts are more (less) accurate than those produced by the driftless random walk model if $DMW_{h,n}^{(i)} > 0$ ($DMW_{h,n}^{(i)} < 0$).

It follows that the null hypothesis of equal MSPEs of our model and the driftless random walk can be expressed as follows.

$$H_0^{(i)}: \text{plim}_n DMW_{h,n}^{(i)} = 0 \text{ for any } i. \quad (6.4)$$

Diebold and Mariano (1995) suggest testing this null using the following asymptotic distribution

$$V_{DMW^{(i)},h}^{-\frac{1}{2}} \sqrt{n_1^{(i)}} DMW_{h,n}^{(i)} \rightarrow_d N(0, 1) \quad (6.5)$$

for $i = 1, \dots, 6$, where $V_{DMW^{(i)},h}$ denotes the LRV estimator of $\sqrt{n_1^{(i)}} DMW_{h,n}^{(i)}$.

Given our forecasting model (6.1), the null tested by the Clark-West Test

is

$$H_0^{cw} : E \left[X_{t,h}^{(i)} \left(e_{t+h}^{(i)} - e_t^{(i)} \right) \right] = 0 \quad (6.6)$$

Clark and West (2006) note that this null is different from the null tested by the DMW test—that the MSPE of the random walk is equal to the MSPE generated by the point forecasts of our model. They show that if the two models are nested, the DMW-test under-rejects the null (6.6). Clark and West (2006) propose a revised statistic—the CW test statistic—that eliminates the bias in favour of the random walk from the DMW-statistic. To see this, decompose the $DMW_{h,n}^{(i)}$ test statistic as follows

$$DMW_{h,n}^{(i)} = -\frac{\widehat{\beta}_{h,n_0}^2}{n_1^{(i)}} \sum_{t=m_0+1}^{n^{(i)}-h} [X_{t,h}^{(i)}]^2 + \frac{2\widehat{\beta}_{h,n_0}}{n_1^{(i)}} \sum_{t=m_0+1}^{n^{(i)}-h} X_{t,h}^{(i)} \left(e_{t+h}^{(i)} - e_t^{(i)} \right),$$

where the first term on the right hand side represents the bias in the DMW test statistic. To eliminate this bias, the CW-test considers the following test statistic

$$CW_{h,n}^{(i)} = DMW_{h,n}^{(i)} + \frac{\widehat{\beta}_{h,n_0}^2}{n_1^{(i)}} \sum_{t=m_0+1}^{n^{(i)}-h} [X_{t,h}^{(i)}]^2 = \frac{2\widehat{\beta}_{h,n_0}}{n_1^{(i)}} \sum_{t=m_0+1}^{n^{(i)}-h} X_{t,h}^{(i)} \left(e_{t+h}^{(i)} - e_t^{(i)} \right).$$

Let $V_{CW^{(i)},h}$ denote the consistent LRV estimator of the CW test statistic. Then using the central limit theorem and the continuous mapping theorem, we have

$$V_{CW,h}^{-\frac{1}{2}} \sqrt{n_1^{(i)}} CW_{h,n}^{(i)} \rightarrow_d N(0, 1). \quad (6.7)$$

As in the recent literature (e.g., Rossi (2013)), we implement the CW test using Newey-West LRV estimator $V_{CW^{(i)},h}$. In addition, we also present test results using the Andrews LRV estimator.

6.1.1 DMW-Test Results

Overall, the DMW-test results show that, over $6m$, $9m$ and $12m$ forecasting horizons, the MSPEs of our point forecasts are significantly lower than the

MSPE of the driftless random walk for the Australian dollar, Canadian dollar, Euro, Yen and the British Pound.

Table 5 contains the $DMW_{h,n}^{(i)}$ test statistics and their p-values for the 6 currencies and the 5 forecasting horizons we consider. The null is rejected if the $DMW_{h,n}^{(i)}$ test statistic is statistically significant and positive. For the one-sided test we consider, a t-value greater than 1.282 implies a 0.1 significance level.

Panel A reports the test results based on the Newey-West LRV estimator. Excluding the Swiss Franc, over the $6m$, $9m$ and $12m$ horizons our point forecasts statistically significantly outperform the driftless random walk in all currencies, except for the British Pound at the $12m$ horizon. At the $3m$ horizon the null is rejected for 3 currencies, while at the $1m$ horizon it is rejected only for the Australian Dollar. Panel B shows that when using the Andrews LRV estimator, the test results are a bit weaker: The $12m$ Yen and the $3m$ British Pound cease to be significant. Still, over $6m$, $9m$ and $12m$ forecasting horizons, the null is rejected in 13 out of 15 currency-horizon pairs.

6.1.2 CW-Test Results

Here, we carry out the standard CW test, in which it is assumed that under the null the exchange rate follows a driftless random walk. Table 6 contains the $CW_{h,n}^{(i)}$ test statistics, and their p-values for the 6 currencies and the 5 horizons. The null is rejected if the $CW_{h,n}^{(i)}$ test statistic is significantly greater than zero. For the one-sided test we consider, a t -value greater than 1.282 implies a 0.1 significance level.

Since, in general, the CW test is more likely to reject the null than the DMW test, we expect that the CW test results will turn out to be more in favor of our point forecasts than those reported in Table 5. Indeed, when we control for auto-correlation using the Newey-West LRV estimator, we can see in Table 6 that the null is rejected in all currencies, except for the Swiss Franc, over the $6m$, $9m$ and $12m$ forecasting horizons. The number of rejections is larger and p-values are lower than the DMW-test results shown in Table 5.

At the $3m(1m)$ horizon the null is rejected in the same three(one) currencies as with the DMW-test. Tests results using the Andrews LRV estimator are similar as those using Newey-West LRV estimator, except for the $12m$ horizon, in which the null is not rejected in two currencies.

7 Conclusion

Exchange rates tend to exhibit swings of appreciation and depreciation. Although these swings can be identified in-sample, they have proven difficult to predict out-of-sample. In this paper, we forecast exchange rates by fitting an autoregressive Markov regime switching model to the speculators' position data in futures markets. The forecasting method we propose combines Engel and Hamilton's (1990) finding that exchanges rate follow long swings with Evans and Lyons' (2004) finding that privately available information about market participants' order flow can predict exchange rates over the short-run.

While Evans and Lyons focus on weekly forecasting horizons and use private information, we concentrate on the 1 -to- 12 months horizons and use public information from the Commitment-of-Traders report. We find that over forecasting horizons ranging from 6 to 12 months, our forecasts outperform those from random walk models for most currencies, except the Swiss Franc. Our directional forecasts have a 58 percent average success ratio and most of our point forecasts have smaller mean-squared-prediction-errors than those implied by the driftless random walk. A battery of econometric tests indicate that, over 6-to-12 months horizons, the outperformance of these two types of forecasts is statistically significant at the 0.1 level for five out of the six most traded currency pairs vis-a-vis the US Dollar: Euro, Japanese Yen, British Pound, Australian Dollar, Canadian Dollar.

In our analysis, we propose a test which weights each directional forecast by the subsequent exchange rate change. Unlike standard binomial tests, this special form gives more weight to the directional forecasts associated with bigger exchange rate moves. Therefore, it is more relevant than the standard binomial test to situations faced by real world investors as it may be linked to

the profitability of some trading strategies. We believe this weighted test may be of interest in other applications.

References

- [1] Alquist, R. and M.D. Chinn (2008): “Conventional and Unconventional Approaches to Exchange Rate Modelling and Assessment,” *International Journal of Finance and Economics* 13, 2-13.
- [2] Andrews, D. (1991): “Heteroskedasticity and Auto-correlation Consistent Covariance Matrix Estimation,” *Econometrica*, 59, 817-854.
- [3] Berkowitz, J. and L. Giorgianni (2001): “Long-Horizon Exchange Rate Predictability?,” *Review of Economics and Statistics* 83(1), 81-91.
- [4] Brunnermeier, M., S. Nagel, and L. Pedersen (2008): “Carry Trades and Currency Crashes,” NBER Chapters, in: *NBER Macroeconomics Annual 2008*, Vol. 23, 313-347.
- [5] Chen, S.L., & Lee, H.Y. (2006): “Why Use Markov-switching Models in Exchange Rate Prediction?,” *Economic Modelling*, 23(4), 662–668.
- [6] Chen Y.C. and K. Rogoff (2003): “Commodity Currencies,” *Journal of International Economics* 60, 133-169.
- [7] Chen, Y.C., K. Rogoff and B. Rossi (2010): “Can Exchange Rates Forecast Commodity Prices?,” *Quarterly Journal of Economics* 125(3), 1145-1194.
- [8] Cheung, Y.W., M.D. Chinn and A.G. Pascual (2005): “Empirical Exchange Rate Models of the Nineties: Are Any Fit to Survive?,” *Journal of International Money and Finance* 24, 1150-1175.
- [9] Chinn, M.D. and R. Meese (1995): “Banking on Currency Forecasts: Is Change in Money Predictable?,” *Journal of International Economics* 38 (1-2), 161-78.

- [10] Clark, T., and K. West (2006): "Using Out-of-Sample Mean Squared Prediction Errors to Test the Martingale Difference Hypothesis," *Journal of Econometrics*, 135, 155-186.
- [11] Clark, T., and K. West (2007): "Approximately Normal Tests for Equal Predictive Accuracy in Nested Models," *Journal of Econometrics*, 138, 291-311.
- [12] Diebold, F., and R. Mariano (1995): "Comparing Predictive Accuracy," *Journal of Business and Economic Statistics*, 13, 253-265.
- [13] Diebold, F., and J. Nason (1990): "Nonparametric Exchange Rate Sized the Importance of Predictable Components in Prediction?," *Journal of International Economics*, 28, 315-335.
- [14] Elliott, G., and T. Ito (1999): "Heterogeneous Expectations and Tests of Efficiency in the Yen/Dollar Forward Exchange Rate Market," *Journal of Monetary Economics*, 43, 435-456.
- [15] Elliott, G., and A. Timmermann (2008): "Economic Forecasts," *Journal of Economic Literature*, 46(1), 3-56.
- [16] Engel, C. (1994): "Can the Markov Switching Model Forecast Exchange rates?," *Journal of International Economics*, 36(1-2), 151-165.
- [17] Engel, C., and J. Hamilton (1990): "Long Swings in the Exchange Rate: Are They in the Data and Do Markets Know It?," *American Economic Review*, 80 (4), 689-713.
- [18] Engel, C., N. Mark and K. West (2007): "Exchange Rate Models Are Not as Bad as You Think," *NBER Macroeconomics Annual*.
- [19] Evans, M., and R. Lyons (2002): "Order Flow and Exchange Rate Dynamics," *Journal of Political Economy* 110, 170-180.

- [20] Faust, J., J.H. Rogers and J.H. Wright (2003): "Exchange Rate Forecasting: the Errors Were Really Made," *Journal of International Economics* 60(1), 35-59.
- [21] Giacomini, R. and H. White (2006): "Tests of Conditional Predictive Ability," *Econometrica*, 74, 1545-1578.
- [22] Gourinchas, P.O. and H. Rey (2007): "International Financial Adjustment," *Journal of Political Economy*, 115(4), 665-703.
- [23] Hamilton, J. (1989): "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica*, 57(2), 357-384.
- [24] Hamilton, J.D., (1990): "Analysis of Time Series Subject to Changes in Regime," *Journal of Econometrics*, 45, 39-70.
- [25] Hansen, B.E. (1992): "Consistent Covariance Matrix Estimation for Dependent Heterogeneous Processes," *Econometrica*, 60, 967-972.
- [26] Hong, H. G. and Yogo, M. (2012): "What Does Futures Market Interest Tell Us about the Macroeconomy and Asset Prices?," *Journal of Financial Economics*, 105(3), 473-490.
- [27] Kilian, L. (1999): "Exchange Rates and Monetary Fundamentals: What Do We Learn From Long-Horizon Regressions?," *Journal of Applied Econometrics* 14(5).
- [28] Kilian, L., and M. Taylor (2003): "Why Is It So Difficult to Beat the Random Walk Forecast of Exchange Rates?," *Journal of International Economics*, 60, 85-107.
- [29] Leitch, G., and Tanner, J. E. (1991), "Econometric Forecast Evaluation: Profits Versus the Conventional Error Measures," *American Economic Review*, 81, 580-590.

- [30] Mark, N.C. (1995): "Exchange Rates and Fundamentals: Evidence on Long-Horizon Predictability," *American Economic Review* 85(1), 201-218.
- [31] Mark, N., and D. Sul (2001): "Nominal Exchange Rates and Monetary Fundamentals: Evidence from a Small Post-Bretton Woods Sample," *Journal of International Economics*, 53, 29-52.
- [32] Meese, R., and K. Rogoff (1983): "Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample," *Journal of International Economics*, 14, 3-24.
- [33] Meese, R., and K. Rogoff(1988): "Was it Real? The Exchange Rate-Interest Differential Relation Over the Modern Floating Rate Period," *Journal of Finance*, 43, 923-948.
- [34] Molodtsova, T. and D.H. Papell (2009): "Out-of-Sample Exchange Rate Predictability with Taylor Rule Fundamentals," *Journal of International Economics*, 77(2), 167-180.
- [35] Moskowitz, T. J, Y. H., and L. H. Pedersen (2011): "Time Series Momentum," *Journal of Financial Economics*, 104(2), 228-250.
- [36] Newey, W.K. and K.D. West (1987): "A Simple, Positive Semi-definite, Heteroscedasticity and Auto-correlation Consistent Covariance Matrix," *Econometrica*, 55, 703-708.
- [37] Newey, W.K. and K.D. West (1994): "Automatic Lag Selection in Covariance Matrix Estimation," *Review of Economic Studies*, 61, 631-653.
- [38] Pesaran M.H. and A., Timermann (1992): "A Simple Nonparametric Test of Predictive Performance," *Journal of Business & Economic Statistics*, 10(4), 461-465.
- [39] Rogoff K. (1996): "The Purchasing Power Parity Puzzle," *Journal of Economic Literature* XXXIV, 647-668.

- [40] Rogoff, K. and V. Stavrageva (2008): “The Continuing Puzzle of Short Horizon Exchange Rate Forecasting,” *NBER Working Paper* No. 14071.
- [41] Rossi, B. (2013): “Exchange Rate Predictability,” mimeo.
- [42] Rossi, B. (2005): “Testing Long-Horizon Predictive Ability with High Persistence, and the Meese-Rogoff Puzzle,” *International Economic Review* 46(1), 61-92.
- [43] Tornell, A. and Yuan, C. (2012): “Speculation and Hedging in the Currency Futures Markets: Are They Informative to the Spot Exchange Rates,” *Journal of Futures Markets*, Vol. 32(2), pp. 122–151.
- [44] West, K. (1996): “Asymptotic Inference About Predictive Ability”, *Econometrica*, 64, 1067-1084.

Table 1. Success Ratio of the Directional Forecasts

Currency	Forecasting Horizon (h)				
	$1m$	$3m$	$6m$	$9m$	$12m$
EUR	46.4%	50.8%	69.0%	66.7%	67.0%
	(125)	(124)	(113)	(99)	(88)
JPY	56.3%	56.1%	60.9%	63.0%	58.6%
	(547)	(574)	(573)	(571)	(560)
GBP	50.6%	55.2%	58.4%	60.3%	60.1%
	(324)	(364)	(365)	(363)	(363)
AUD	58.0%	67.4%	72.4%	72.7%	77.7%
	(269)	(282)	(283)	(275)	(265)
CAD	54.9%	55.3%	56.0%	68.0%	66.8%
	(381)	(398)	(398)	(391)	(391)
CHF	52.1%	49.1%	51.4%	50.2%	50.1%
	(576)	(621)	(623)	(614)	(603)

Notes: 1. The forecast success ratio is the ratio of the number of successful depreciation and appreciation forecasts divided by the total number of depreciation and appreciation forecasts. 2. The total number of appreciation and depreciation forecasts is in parentheses. 3. Our sample starts on 10/02/1992 for all currencies except the Euro, which starts on 01/08/1999. Our sample ends on 02/08/2013 for all currencies. 4. For each currency, our directional forecasts start the week after the first estimation of the MSM, and end h weeks prior to the end of the COT sample. Thus, our directional forecasts start on 01/20/1995 for AUD and GBP; 09/01/1994 for CAD; 04/15/1994 for CHF and JPY; and 05/04/2001 for the Euro. 4. The dates of the last forecasts are 01/18/2013 for $h = 1m$, 11/16/2012 for $h = 3m$, 08/24/2012 for $h = 6m$, 05/25/2012 for $h = 9m$, and 03/01/2012 for $h = 12m$.

Table 2. Success Ratio (Two Subperiods) of the Directional Forecasts

Currency	Forecasting Horizon (h)				
	1m	3m	6m	9m	12m
First half period: 04/15/1994 - 11/29/2002					
EUR	28.6%	57.1%	71.4%	28.6%	57.1%
	(7)	(7)	(7)	(7)	(7)
JPY	59.5%	57.4%	64.4%	66.0%	60.4%
	(232)	(242)	(247)	(247)	(240)
GBP	51.8%	54.5%	58.3%	63.1%	63.5%
	(218)	(244)	(242)	(241)	(241)
AUD	60.0%	80.0%	90.0%	87.5%	75.0%
	(10)	(10)	(10)	(8)	(4)
CAD	47.2%	45.2%	30.8%	33.8%	28.8%
	(89)	(84)	(78)	(77)	(66)
CHF	54.4%	50.2%	48.8%	47.7%	52.7%
	(307)	(307)	(297)	(283)	(273)
Second half period: 12/06/2002 - 02/08/2013					
EUR	47.5%	50.4%	68.9%	69.6%	67.9%
	(118)	(117)	(106)	(92)	(81)
JPY	54.0%	55.1%	58.3%	60.8%	57.2%
	(315)	(332)	(326)	(324)	(320)
GBP	48.1%	56.7%	58.5%	54.9%	53.3%
	(106)	(120)	(123)	(122)	(122)
AUD	57.9%	66.9%	71.8%	72.3%	77.8%
	(259)	(272)	(273)	(267)	(261)
CAD	57.2%	58.0%	62.2%	76.4%	74.5%
	(292)	(314)	(320)	(314)	(325)
CHF	49.4%	48.1%	53.7%	52.3%	47.9%
	(269)	(314)	(326)	(331)	(330)

Note: The information on our directional forecasts is described in the notes to Table 1.

Table 3: Weighted Directional Forecasts Test

Currency	Forecasting Horizon (h)				
	1m	3m	6m	9m	12m
Panel A: Newey-West					
EUR	-.0003 (-0.416)	.0012 (0.779)	.0034** (1.853)	.0051*** (2.568)	.0044** (2.320)
JPY	.0020** (1.764)	.0024 (0.788)	.0071** (1.674)	.0115*** (2.376)	.0097* (1.537)
GBP	.0002 (0.247)	.0023** (1.686)	.0040** (2.149)	.0057*** (2.534)	.0061** (2.054)
AUD	.0018** (1.739)	.0038* (1.506)	.0105*** (2.743)	.0163*** (3.593)	.0187*** (3.321)
CAD	.0008 (1.219)	.0030** (1.803)	.0037* (1.477)	.0109*** (4.064)	.0156*** (4.480)
CHF	.0014 (1.154)	-.0012 (-0.404)	-.0023 (-0.566)	.0016 (0.328)	-.0034 (-0.567)
Panel B: Andrews					
EUR	-.0003 (-0.417)	.0012 (0.685)	.0034** (1.687)	.0051*** (2.568)	.0044*** (2.334)
JPY	.0020** (1.654)	.0024 (0.608)	.0071 (1.187)	.0115* (1.620)	.0097 (0.993)
GBP	.0002 (1.241)	.0023* (1.476)	.0040* (1.643)	.0057** (1.806)	.0061* (1.569)
AUD	.0018** (1.705)	.0038 (1.268)	.0105** (1.942)	.0163*** (2.412)	.0187** (2.103)
CAD	.0008 (1.191)	.0030*** (1.509)	.0037 (1.164)	.0109*** (3.139)	.0156*** (3.035)
CHF	.0014* (1.375)	-.0012 (-0.392)	-.0023 (-0.426)	.0016 (0.242)	-.0034 (-0.400)

Note: 1. Panel A and B report the test results using Newey-West and Andrews LRV estimators to control for auto-correlation, respectively. 2. t-values are in parentheses. We use the test as an one-sided test. Critical values from a standard normal distribution are used for inference. *, ** and *** indicate significance at 10, 5 and 1 percent, respectively. 3. The information on our directional forecasts is described in the notes to Table 1.

Table 4: Binomial Directional Test

Currency	Forecasting Horizon (h)				
	1m	3m	6m	9m	12m
Panel A: Newey-West					
EUR	-.0033 (-0.129)	.0284 (0.896)	.0963*** (2.921)	.0775*** (2.942)	.0658*** (2.770)
JPY	.0684** (2.063)	.0722* (1.578)	.1310*** (2.594)	.1606*** (2.916)	.1082** (1.859)
GBP	.0017 (0.065)	.0405 (1.047)	.0597* (1.539)	.0763** (1.775)	.0774** (1.723)
AUD	.0291 (1.078)	.0695** (2.098)	.0966*** (2.661)	.0896*** (2.468)	.1102*** (2.964)
CAD	.0258 (0.877)	.0334 (0.823)	.0252 (0.556)	.1061** (2.336)	.1028** (2.276)
CHF	.0269 (0.746)	-.0011 (-0.022)	.0319 (0.566)	.0295 (0.5012)	.0339 (0.565)
Panel B: Andrews					
EUR	-.0033 (-0.127)	.0284 (0.801)	.0963*** (2.403)	.0775*** (2.880)	.0658*** (2.728)
JPY	.0684** (2.045)	.0722* (1.368)	.1310** (1.998)	.1606** (2.086)	.1082 (1.287)
GBP	0.002 (0.064)	0.041 (0.946)	0.060* (1.336)	0.076* (1.402)	0.077* (1.353)
AUD	.0291 (1.091)	.0695** (1.867)	.0966** (2.092)	.0896** (1.916)	.1102** (2.037)
CAD	.0258 (0.866)	.0334 (0.723)	.0252 (0.447)	.1061** (1.790)	.1028** (1.691)
CHF	.0269 (0.781)	-.0011 (-0.022)	.0319 (0.442)	.0295 (0.374)	.0339 (0.404)

Note: 1. Panel A and B report the test results using Newey-West and Andrews LRV estimators to control for auto-correlation, respectively. 2. t-values are in parentheses. We use the test as an one-sided test. Critical values from a standard normal distribution are used for inference. *, ** and *** indicate significance at 10, 5 and 1 percent, respectively. 3. The information on our directional forecasts is described in the notes to Table 1.

Table 5: Diebold-Mariano-West Test

Currency	Forecasting Horizon (h)				
	1m	3m	6m	9m	12m
Panel A: Newey-West					
EUR	-0.287 (0.613)	1.111 (0.133)	1.614* (0.053)	2.541*** (0.006)	2.384*** (0.009)
JPY	0.770 (0.221)	1.057 (0.145)	2.118** (0.017)	2.261** (0.012)	1.438* (0.075)
GBP	0.003 (0.499)	1.554* (0.060)	2.251** (0.012)	2.065** (0.019)	1.089 (0.138)
AUD	1.622* (0.052)	1.964** (0.025)	2.995*** (0.001)	4.020*** (0.000)	3.509*** (0.000)
CAD	0.643 (0.260)	1.675** (0.047)	1.819** (0.034)	3.630*** (0.000)	4.787*** (0.000)
CHF	-0.102 (0.541)	-0.678 (0.751)	-1.260 (0.896)	-0.829 (0.797)	-1.060 (0.856)
Panel B: Andrews					
EUR	-0.286 (0.613)	0.979 (0.164)	1.356* (0.088)	2.473*** (0.007)	2.384*** (0.009)
JPY	0.685 (0.247)	0.793 (0.214)	1.408* (0.080)	1.337* (0.091)	0.794 (0.214)
GBP	0.003 (0.499)	1.249 (0.106)	1.548* (0.061)	1.286* (0.099)	0.728 (0.233)
AUD	1.586* (0.056)	1.596* (0.055)	2.038** (0.021)	2.531*** (0.006)	2.044** (0.020)
CAD	0.650 (0.258)	1.366* (0.086)	1.370* (0.085)	2.524*** (0.006)	2.917*** (0.002)
CHF	-0.123 (0.549)	-0.693 (0.756)	-0.845 (0.801)	-0.549 (0.708)	-0.673 (0.750)

Note: 1. Panel A and B report the test results using the Newey-West and Andrews LRV estimators respectively. 2. p-values are in parentheses. We use the test as an one-sided test. Critical values from a standard normal distribution are used for inference. *, ** and *** indicate significance at 10, 5 and 1 percent, respectively. 3. The first point forecast starts on 04/09/1999 for five currencies excluding the Euro. For the Euro it starts on 05/04/2001. The dates of the last point forecasts are 01/18/2013 for $h = 1m$, 11/16/2012 for $h = 3m$, 08/24/2012 for $h = 6m$, 05/25/2012 for $h = 9m$, and 03/01/2012 for $h = 12m$.

Table 6. Clark-West Test

Currency	Forecasting Horizon (h)				
	1m	3m	6m	9m	12m
Panel A: Newey-West					
EUR	-0.001 (0.500)	1.173 (0.120)	1.872** (0.031)	2.767*** (0.003)	2.547*** (0.005)
JPY	1.221 (0.111)	1.157 (0.124)	2.521*** (0.006)	2.713*** (0.003)	1.696** (0.045)
GBP	0.378 (0.353)	1.643** (0.050)	2.540*** (0.006)	2.420*** (0.008)	1.305* (0.096)
AUD	1.884** (0.030)	2.028** (0.021)	3.234*** (0.001)	4.267*** (0.000)	3.650*** (0.000)
CAD	1.126 (0.130)	1.777** (0.038)	2.200** (0.014)	4.113*** (0.000)	5.059*** (0.000)
CHF	0.327 (0.372)	-0.584 (0.720)	-0.886 (0.812)	-0.397 (0.654)	-0.828 (0.796)
Panel B: Andrews					
EUR	-0.001 (0.500)	1.033 (0.151)	1.568* (0.058)	2.647*** (0.004)	2.516*** (0.006)
JPY	1.084 (0.139)	0.867 (0.193)	1.663** (0.048)	1.596* (0.055)	0.935 (0.175)
GBP	0.380 (0.352)	1.314* (0.094)	1.711** (0.044)	1.480* (0.069)	0.870 (0.192)
AUD	1.838** (0.033)	1.644** (0.050)	2.191** (0.014)	2.647*** (0.004)	2.114** (0.017)
CAD	1.121 (0.131)	1.445* (0.074)	1.646** (0.050)	2.821*** (0.002)	3.043*** (0.001)
CHF	0.390 (0.348)	-0.597 (0.725)	-0.592 (0.723)	-0.261 (0.603)	-0.520 (0.699)

Note: 1. Panel A and B report the test results using Newey-West and Andrews LRV estimators to control for auto-correlation. 2. p-values are in parentheses. We use the test as an one-sided test. Critical values from a standard normal distribution are used for inference. *, ** and *** indicate significance at 10, 5 and 1 percent, respectively. 3. The information on our point forecasts is described in the notes to Table 5.

Appendix For Online Publication

A Equilibrium of the Dynamic Model (3.11)-(3.12)

In this appendix, we show that the dynamic model (3.11)-(3.12) has a convergent fixed point. From the equations in (3.7), we can write

$$\sigma_t^2 = \frac{\sigma_{t-1}^2 + \sigma_u^2}{\sigma_{t-1}^2 + \sigma_u^2 + \sigma_v^2} \sigma_v^2$$

which implies that

$$\begin{aligned} \sigma_{t+1}^2 - \sigma_t^2 &= \left[\frac{\sigma_t^2 + \sigma_u^2}{\sigma_t^2 + \sigma_u^2 + \sigma_v^2} - \frac{\sigma_{t-1}^2 + \sigma_u^2}{\sigma_{t-1}^2 + \sigma_u^2 + \sigma_v^2} \right] \sigma_v^2 \\ &= \frac{(\sigma_t^2 - \sigma_{t-1}^2) \sigma_v^2}{(\sigma_t^2 + \sigma_u^2 + \sigma_v^2)(\sigma_{t-1}^2 + \sigma_u^2 + \sigma_v^2)} \sigma_v^2. \end{aligned}$$

This means that if we have $\sigma_t^2 = \sigma_{t-1}^2$ for some $t = t^*$, then $\sigma_{t+1}^2 = \sigma_t^2$ for any $t \geq t^*$. Otherwise, we have

$$\frac{\sigma_{t+1}^2 - \sigma_t^2}{\sigma_t^2 - \sigma_{t-1}^2} = \frac{\sigma_v^4}{(\sigma_t^2 + \sigma_u^2 + \sigma_v^2)(\sigma_{t-1}^2 + \sigma_u^2 + \sigma_v^2)}$$

which together with

$$(\sigma_t^2 + \sigma_u^2 + \sigma_v^2)(\sigma_{t-1}^2 + \sigma_u^2 + \sigma_v^2) > \sigma_v^4$$

implies that

$$\frac{\sigma_{t+1}^2 - \sigma_t^2}{\sigma_t^2 - \sigma_{t-1}^2} < 1.$$

Hence, by the contraction mapping theorem we know that σ_t^2 converges to a unique equilibrium (fixed point). Moreover, by definition,

$$\theta_t = \frac{\sigma_{t-1}^2 + \sigma_\varepsilon^2}{\sigma_t^2 + \sigma_\varepsilon^2} (1 - k_t) = \frac{\sigma_{t-1}^2 + \sigma_\varepsilon^2}{\sigma_t^2 + \sigma_\varepsilon^2} \frac{\sigma_v^2}{\sigma_t^2 + \sigma_u^2 + \sigma_v^2}$$

which together with the convergence of σ_t^2 implies that the limit of θ_t is between 0 and $\sigma_v^2(\sigma_u^2 + \sigma_v^2)^{-1}$.

B Constructing the LRV Estimators

For any weakly dependent process $\{W_{t,n}\}_{t=1}^n$ with

$$E[W_{t,n}] = 0 \text{ for all } t \text{ and } n, \quad (\text{B.1})$$

and finite positive LRV V_W , its sample autocovariance can be defined as

$$\Gamma_{W,n}(j) = \frac{1}{n-j} \sum_{t=1}^{n-j} (W_{t,n} - \bar{W}_n) (W_{t+j,n} - \bar{W}_n) \quad (\text{B.2})$$

for $j = 0, \dots, n-1$. It is clear that the sample autocovariance satisfies $\Gamma_{W,n}(-j) = \Gamma_{W,n}(j)$ for $j = 0, \dots, n-1$. Note that the sample autocovariance is sample mean centered, which improves the power of the test of the hypothesis in (B.1).

The kernel based LRV estimator for $\{W_{t,n}\}_{t=1}^n$ is then defined as

$$V_{W,n} = \sum_{j=-n+1}^{n+1} K(j/M) \Gamma_{W,n}(j) \quad (\text{B.3})$$

where $K(\cdot)$ is some kernel smoothing function with bandwidth M . Under some regularity conditions (see, e.g., Newey and West (1987), Andrews (1991) and Hansen (1992)), there is

$$V_{W,n} \rightarrow_p V_W. \quad (\text{B.4})$$

One key condition for the above consistency result is that M goes to infinity at certain rate. In finite samples, there are two different rules of selecting M : One is the rule proposed in Newey and West (1994) and the other is the parametric (AR(1)) approximation rule in Andrews (1991).

In the rest of this appendix, we briefly describe how to construct the LRV estimators for the test statistics presented in the main text. The Newey-West

and Andrews LRV estimators can be constructed using the formula in (B.2). Hence for each test statistic, we only need to define its corresponding " $W_{t,n}$ " for the construction of LRV estimators, which are summarized in Table B.1. For the ease of notation, we ignore the index " i " in each test statistic.

Table B.1 Construction of LRV Estimators

Test Statistic	$W_{t,n}$
$T_{a,n}$	$D_{t,h}(e_{t+h} - e_t)$
$T_{b,n}$	$(D_{t,h} - \bar{D}_{n,h})(R_{t,h} - \bar{R}_{n,h})$
$T_{c,n}$	$D_{t,h}(e_{t+h} - e_t - \bar{e}_{n,h})$
$T_{d,n}$	$(D_{t,h} - \bar{D}_{n,h})(\hat{R}_{t,h} - \hat{\bar{R}}_{n,h})$
$CM_{h,n}$	$2\hat{\beta}_{h,n_0} X_{t,h} (e_{t+h} - e_t)$
$DMW_{h,n}^{dl}$	$(e_{t+h} - e_t)^2 - (e_{t+h} - e_t - X_{t,h}\hat{\beta}_{h,m_0})^2$
$DMW_{h,n}^d$	$(e_{t+h} - e_t - \hat{\eta}_{h,m_0})^2 - (e_{t+h} - e_t - X_{t,h}\hat{\beta}_{h,m_0})^2$

Table B.1: Note: $DM_{h,n}^{dl}$ and $DM_{h,n}^d$ refer to the DM test statistics against the random walk with and without the drift term respectively. Similarly, $CW_{h,n}^{dl}$ denotes the CW test statistic

C Comparison with the Random Walk with Drift

In this section, we compare our directional and point forecasts with those generated by the random walk model with drift:

$$e_{t+1}^{(i)} = e_t^{(i)} + c^{(i)} + \varepsilon_{t+1}^{(i)}, \quad (\text{C.1})$$

where $\{\varepsilon_t^{(i)}\}$ is a white noise process with mean zero and variance $\sigma_{i,\varepsilon}^2$, and $c^{(i)}$ is some finite constant. As the tests results below show, driftless random walk is a tougher benchmark to beat than the random walk with drift. The motivation is that we want to evaluate the robustness of our empirical findings by considering different random walk models.

C.1 Directional Tests

When the null hypothesis is a random walk with drift, the optimal h -period ahead forecast of $e_{t+h}^{(i)}$ is $hc^{(i)} + e_t^{(i)}$ given the martingale assumption on $\{\varepsilon_t^{(i)}\}$. Because the term $hc^{(i)}$ is unknown and can be estimated by $\bar{e}_{n,h}^{(i)} = n_1^{-1} \sum_{t=m_0}^{n-h} (e_{t+h}^{(i)} - e_t^{(i)})$, the feasible point forecast from a random walk with drift is $e_t^{(i)} + \bar{e}_{n,h}^{(i)}$. Thus, we consider the following weighted directional forecast evaluation test statistic

$$T_{c,n}^{(i)} = \frac{1}{n_1^{(i)}} \sum_{t=n_0^{(i)}}^{n-h} D_{t,h}^{(i)} (e_{t+h}^{(i)} - e_t^{(i)} - \bar{e}_{n,h}^{(i)}), \quad (\text{C.2})$$

where $D_{t,h}^{(i)}$ is defined in (5.4). The null hypothesis is then

$$H_0 : \quad E \left[D_{t,h}^{(i)} (e_{t+h}^{(i)} - e_t^{(i)} - hc^{(i)}) \right] = 0 \quad (\text{C.3})$$

for any pre-specified h and any t , which means that after the adjustment of the deterministic trend, our directional forecasts are uncorrelated with future realized exchange rate changes. Let $V_{T_{c,n}^{(i)}}$ denote the consistent estimator of $T_{c,n}^{(i)}$. The inference of the null in (C.3) is then based on the following asymptotic theory

$$V_{T_{c,n}^{(i)}}^{-\frac{1}{2}} \sqrt{n_1^{(i)}} T_{c,n}^{(i)} \rightarrow_d N(0, 1). \quad (\text{C.4})$$

Table C.1 presents the values of the $T_{c,n}^{(i)}$ statistic, and its t -values, for the 6 currencies and the 5 horizons we consider. These tables are the counterparts of Tables 3, in the sense that they use the same variance estimators. The only difference is that they test the null that our directional forecast is uncorrelated with future realized exchange rate changes adjusted by the deterministic trend (i.e., $E[D_{t,h}^{(i)}(e_{t+h}^{(i)} - e_t^{(i)} - hc^{(i)})] = 0$). The overall results confirm that our directional forecasts provide strong evidence of exchange rate predictability over the same forecasting horizons (i.e., $6m$, $9m$ and $12m$) and they are robust after controlling the deterministic trend.

Table C.1: Weighted Directional Test Against the Drift Random Walk Model

Currency	Forecasting Horizon (h)				
	1m	3m	6m	9m	12m
Panel A: Newey-West					
EUR	-.0001 (-0.028)	.0021* (1.368)	.0048*** (2.409)	.0066*** (2.866)	.0057*** (2.508)
JPY	.0020** (1.766)	.0024 (0.786)	.0072** (1.699)	.0119** (2.444)	.0102* (1.619)
GBP	.0002 (0.247)	.0023** (1.685)	.0040** (2.144)	.0057*** (2.532)	.0061** (2.044)
AUD	.0015* (1.451)	.0025 (1.006)	.0079** (2.120)	.0124*** (2.905)	.0138*** (2.578)
CAD	.0006 (0.858)	.0019 (1.213)	.0013 (0.569)	.0070** (2.848)	.0105** (3.388)
CHF	.0017* (1.364)	-.0003 (-0.098)	-.0004 (-0.109)	.0042 (0.915)	-.0002 (-0.041)
Panel B: Andrews					
EUR	-.0001 (-0.0279)	.0021 (1.167)	.0048** (2.088)	.0066*** (2.757)	.0057*** (2.488)
JPY	.0020** (1.654)	.0024 (0.604)	.0072 (1.196)	.0119* (1.637)	.0102 (1.031)
GBP	.0002 (0.241)	.0023* (1.475)	.0040* (1.639)	.0057** (1.804)	.0061* (1.561)
AUD	.0015* (1.427)	.0025 (0.860)	.0079* (1.5178)	.0124** (2.000)	.0138** (1.682)
CAD	.0006 (0.849)	.0019 (1.034)	.0013 (0.452)	.0070** (2.255)	.0105*** (2.350)
CHF	.0017* (1.605)	-.0003 (-0.096)	-.0004 (-0.084)	.0042 (0.695)	-.0002 (-0.030)

Note: 1. Panel A and B report the test results using Newey-West and Andrews LRV estimators to control for auto-correlation, respectively. 2. t-values are in parentheses. We use the test as an one-sided test. Critical values from a standard normal distribution are used for inference. *, ** and *** indicate significance at 10, 5 and 1 percent, respectively. 3. The information on our directional forecasts is described in the notes to Table 1.

Table C.2: Binomial Directional Test Against Drift Random Walk Model

Currency	Forecasting Horizon (h)				
	$1m$	$3m$	$6m$	$9m$	$12m$
Panel A: Newey-West					
EUR	.0070 (0.260)	.0300 (0.982)	.0802*** (2.466)	.0911*** (3.311)	.0623*** (2.655)
JPY	.0684** (2.046)	.0764* (1.632)	.1320*** (2.625)	.1530*** (2.772)	.0914* (1.577)
GBP	.0017 (0.065)	.0405 (1.047)	.0596* (1.536)	.0761** (1.771)	.0772** (1.720)
AUD	.0289 (1.066)	.0618** (1.960)	.0908*** (2.499)	.0793** (2.139)	.1091*** (2.955)
CAD	.0183 (0.638)	.0444 (1.076)	.0330 (0.722)	.1130*** (2.579)	.1333*** (2.872)
CHF	.0232 (0.635)	.0048 (0.097)	.0147 (0.260)	.0812* (1.425)	.0963* (1.593)
Panel B: Andrews					
EUR	.0070 (0.254)	.0300 (0.903)	.0802** (2.027)	.0911*** (3.001)	.0623*** (2.591)
JPY	.0684** (2.024)	.0764* (1.386)	.1320** (2.048)	.1530** (1.976)	.0914 (1.084)
GBP	.0017 (0.0638)	.0405 (0.946)	.0596* (1.334)	.0761* (1.399)	.0772* (1.350)
AUD	.0289 (1.083)	.0618** (1.847)	.0908** (1.968)	.0793** (1.681)	.1091** (2.067)
CAD	.0183 (0.639)	.0444 (0.937)	.0330 (0.580)	.1130** (2.089)	.1333** (2.128)
CHF	.0232 (0.662)	.0048 (0.097)	.0147 (0.199)	.0812 (1.101)	.0963 (1.154)

Note: 1. Panel A and B report the test results using Newey-West and Andrews LRV estimators to control for auto-correlation, respectively. 2. t-values are in parentheses. We use the test as an one-sided test. Critical values from a standard normal distribution are used for inference. *, ** and *** indicate significance at 10, 5 and 1 percent, respectively. 3. The information on our directional forecasts is described in the notes to Table 1.

For the Binomial directional test, the test statistic is

$$T_{d,n}^{(i)} = \frac{1}{n_1^{(i)}} \sum_{t=n_0^{(i)}}^{n-h} D_{t,h}^{(i)} \widehat{R}_{t,h}^{(i)} - \frac{1}{n_1^{(i)}} \sum_{t=n_0^{(i)}}^{n-h} D_{t,h}^{(i)} \frac{1}{n_1^{(i)}} \sum_{t=n_0^{(i)}}^{n-h} \widehat{R}_{t,h}^{(i)}$$

where $\widehat{R}_{t,h}^{(i)} = 1$ if $e_{t+h}^{(i)} - e_t^{(i)} - h\bar{e}_n^{(i)} \geq 0$, and $\widehat{R}_{t,h}^{(i)} = -1$ otherwise, $\bar{e}_n^{(i)} = \frac{1}{n^{(i)}} \sum_{t=1}^{n^{(i)}} e_t^{(i)}$. Let $V_{T_{d,n}^{(i)}}$ denote the consistent LRV estimator of $T_{d,n}^{(i)}$. Then we have

$$V_{T_{d,n}^{(i)}}^{-\frac{1}{2}} \sqrt{n_1^{(i)}} T_{d,n}^{(i)} \rightarrow_d N(0, 1) ,$$

which is used in testing the following null hypothesis

$$H_0: Cov(D_{t,h}^{(i)}, \bar{R}_{t,h}^{(i)}) = 0, \quad (C.5)$$

where $\bar{R}_{t,h}^{(i)} = 1$ if $e_{t+h}^{(i)} - e_t^{(i)} - hc^{(i)} \geq 0$, and $\bar{R}_{t,h}^{(i)} = -1$ otherwise.

Tables C.2 presents the values of the $T_{d,n}^*$ statistic, and its t -values, for the 6 currencies and the 5 horizons we consider. They are analogous to Table 4 in the sense that they use the same variance estimators. The only difference is that they test the null that our directional forecast is uncorrelated with the directional sign of future realized exchange rate changes adjusted by the deterministic trend (i.e., $Cov(D_{t,h}^{(i)}, \bar{R}_{t,h}^{(i)}) = 0$). The overall results confirm that the previous results remain valid even if the realized exchange rate trend is controlled in the determination of the sign of change. Furthermore, some evidence of predictability is found for the Swiss Franc at $9m$ and $12m$ horizons.

C.2 The DMW Test

In this subsection, we compare the point prediction accuracy of our model with the random walk with a drift term, i.e.

$$e_{t+h}^{(i)} = \eta_h + e_t^{(i)} + \varepsilon_{t+h}^{(i)} \quad (C.6)$$

Table C.3: DMW Test Against Drift Random Walk Model

Currency	Forecasting Horizon (h)				
	$1m$	$3m$	$6m$	$9m$	$12m$
Panel A: Newey-West					
EUR	2.202** (0.014)	2.300** (0.011)	3.021*** (0.001)	3.597*** (0.000)	4.258*** (0.000)
JPY	1.744** (0.041)	1.522* (0.064)	2.964*** (0.002)	4.127*** (0.000)	4.244*** (0.000)
GBP	0.411 (0.341)	0.524 (0.300)	0.931 (0.176)	1.051 (0.147)	1.251 (0.105)
AUD	2.842*** (0.002)	1.841** (0.033)	2.576*** (0.005)	3.305*** (0.000)	3.817*** (0.000)
CAD	2.482*** (0.007)	2.543*** (0.005)	3.208*** (0.001)	4.140*** (0.000)	4.959*** (0.000)
CHF	2.088** (0.018)	2.413*** (0.008)	3.038*** (0.001)	4.293*** (0.000)	5.666*** (0.000)
Panel B: Andrews					
EUR	1.213 (0.113)	1.813** (0.035)	1.985** (0.024)	1.914** (0.028)	1.901** (0.029)
JPY	0.901 (0.184)	1.222 (0.111)	1.846** (0.032)	1.833** (0.033)	1.585* (0.056)
GBP	0.212 (0.416)	0.373 (0.355)	0.551 (0.291)	0.516 (0.303)	0.537 (0.295)
AUD	1.528* (0.063)	1.364* (0.086)	1.690** (0.046)	1.875** (0.030)	1.921** (0.027)
CAD	1.432* (0.076)	1.818** (0.035)	2.159** (0.015)	2.379*** (0.009)	2.583*** (0.005)
CHF	1.287* (0.099)	1.865** (0.031)	1.946** (0.026)	2.339*** (0.010)	2.898*** (0.002)

Note: 1. Panel A and B report the test results using Newey-West and Andrews LRV estimators to control for auto-correlation, respectively. 2. t-values are in parentheses. We use the test as an one-sided test. Critical values from a standard normal distribution are used for inference. *, ** and *** indicate significance at 10, 5 and 1 percent, respectively. 3. The information on our point forecasts is described in the notes to Table 5.

where η_h is a finite constant. The h -period ahead prediction based on the random walk model is

$$\bar{e}_{t+h}^{(i)} = \hat{\eta}_{h,m_0} + e_t^{(i)} \text{ for } m_0 + 1 \leq t \leq n$$

where

$$\widehat{\eta}_{h,m_0} = \frac{\sum_{i \in I} \sum_{t=1}^{m_0-h} (e_{t+h}^{(i)} - e_t^{(i)})}{5(m_0 - h)}, \quad I = \{ad, bp, cd, jy, sf\}.$$

Using the quadratic loss function, we can evaluate the accuracy of our point forecast by the statistic $DMW_{h,n}^{(i)}$, i.e.

$$DMW_{h,n}^{(i)} = \frac{1}{n_1^{(i)}} \sum_{t=m_0+1}^{n^{(i)}-h} \left[\left(\widehat{\eta}_{h,m_0} + e_t^{(i)} - e_{t+h}^{(i)} \right)^2 - \left(\widehat{\beta}_{h,m_0} X_{t,h}^{(i)} + e_t^{(i)} - e_{t+h}^{(i)} \right)^2 \right] \quad (\text{C.7})$$

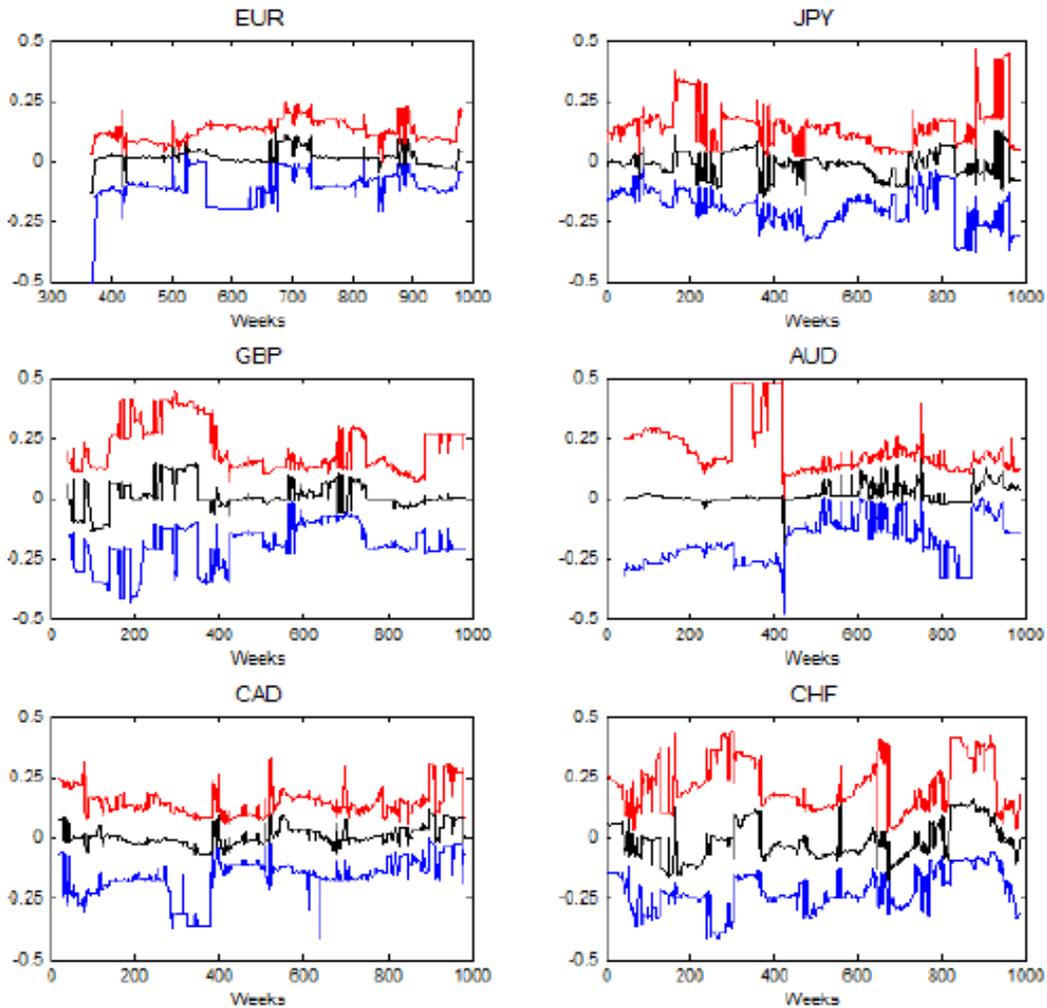
where $\widehat{\beta}_{h,m_0}$ is defined in (6.2).

From the definition of the $DMW_{h,n}^{(i)}$ statistic in (C.7), we can construct consistent LRV estimators. Hence, asymptotic theory similar to that stated in (6.5) can be used to test null hypothesis (6.4).

Table C.3 contains the $DMW_{h,n}^{(i)}$ test statistics and their p-values for the 6 currencies and the 5 horizons we consider. This table tests the null that the MSPE of a random walk with drift is equal to that generated by of our point forecasts, by using the Newey-West LRV estimator and the Andrews LRV estimator. The null is rejected if the $DMW_{h,n}^{(i)}$ test statistic is significantly greater than zero. For the one-sided test we consider, a t -value greater than 1.282, which implies a 0.1 significance level. The overall results show that our point forecasts significantly outperform the random walk with drift in 5 currencies for most forecasting horizons.

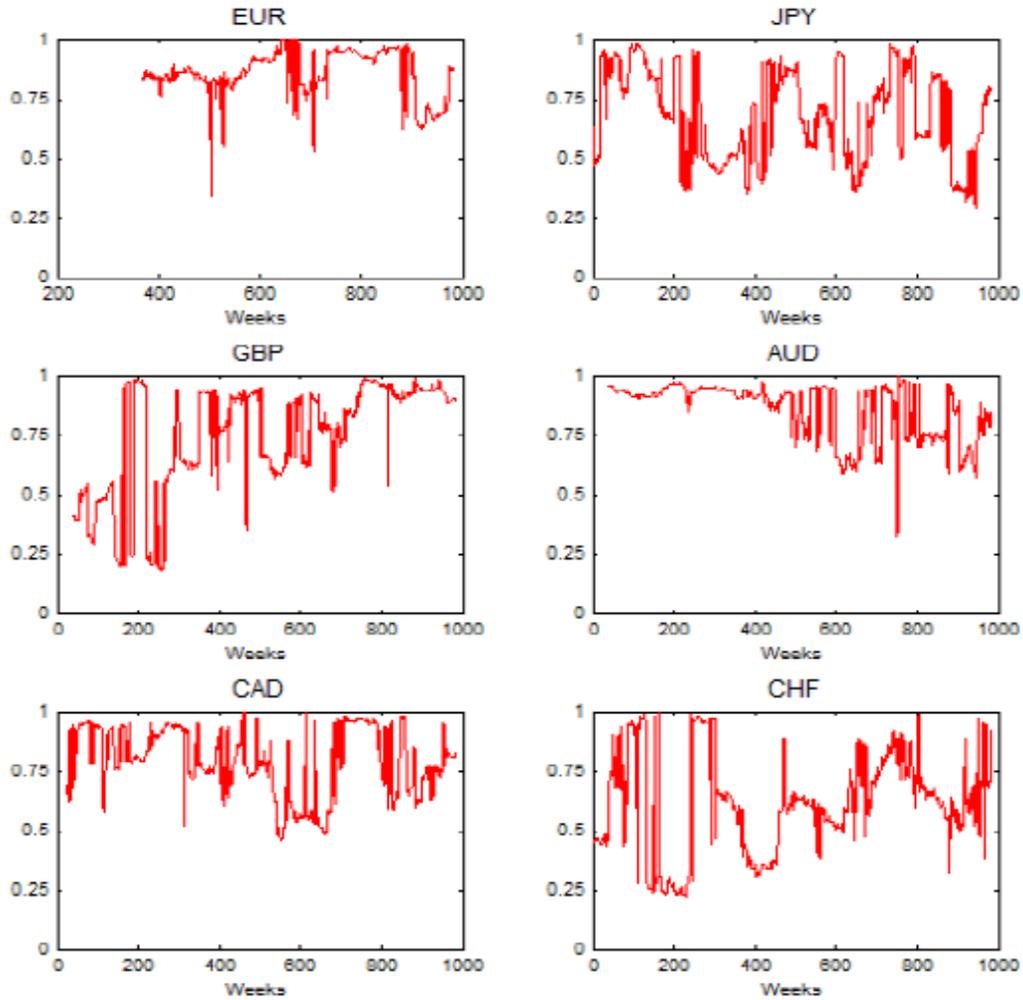
D Figures

Figure D.1. Evolution of μ Estimators



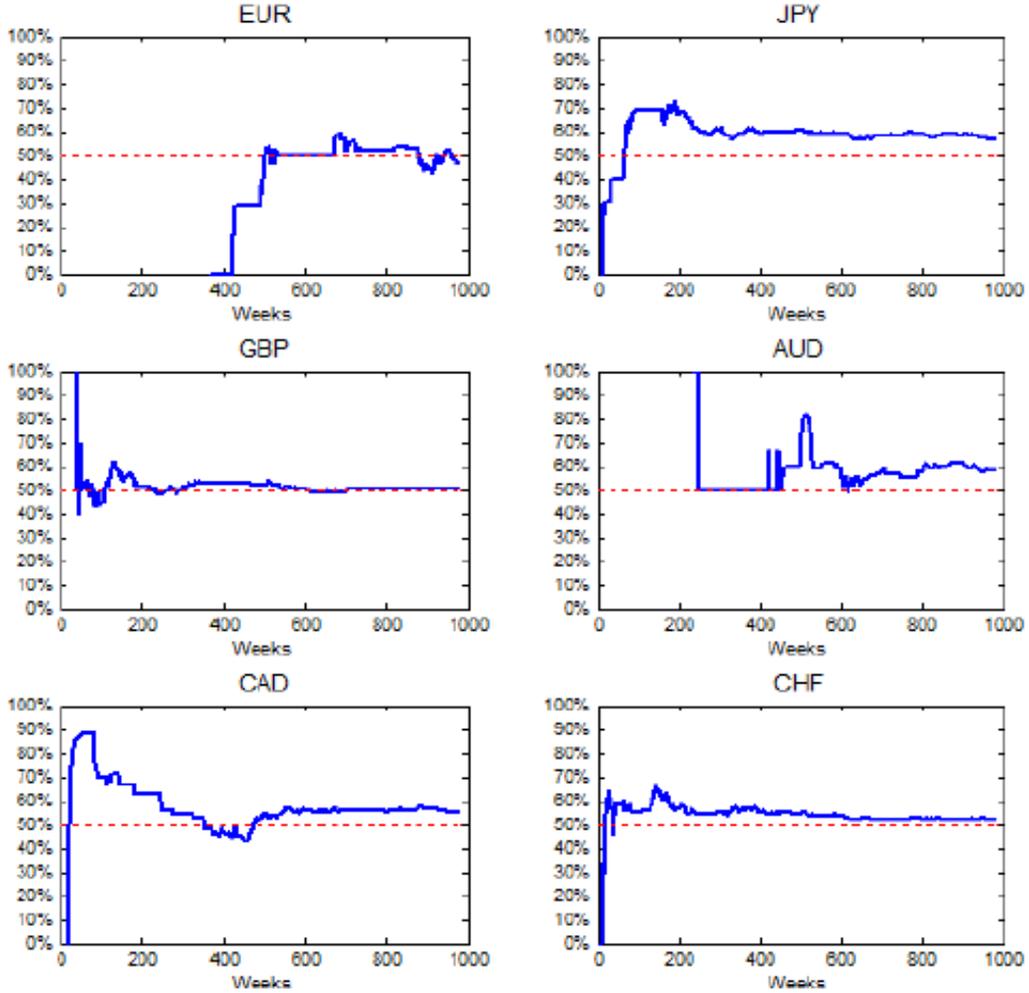
Notes: 1. This graph presents the estimated means of each state in the AR(1) MSM. The red solid line (blue line) depicts the estimated mean of the up (down) state, while the black line describes the estimated mean of the range state. 2. The rolling window size in the estimation of the MSM is 120 weeks for EUR, GBP and AUD, 100 weeks for CAD, and 80 weeks for JPY and CHF. The estimation begins on 01/20/1995 for GBP and AUD, on 09/01/1994 for CAD, on 04/15/1994 for JPY and CHF, and on 05/04/2001 for EUR.

Figure D.2. Evolution of θ Estimators



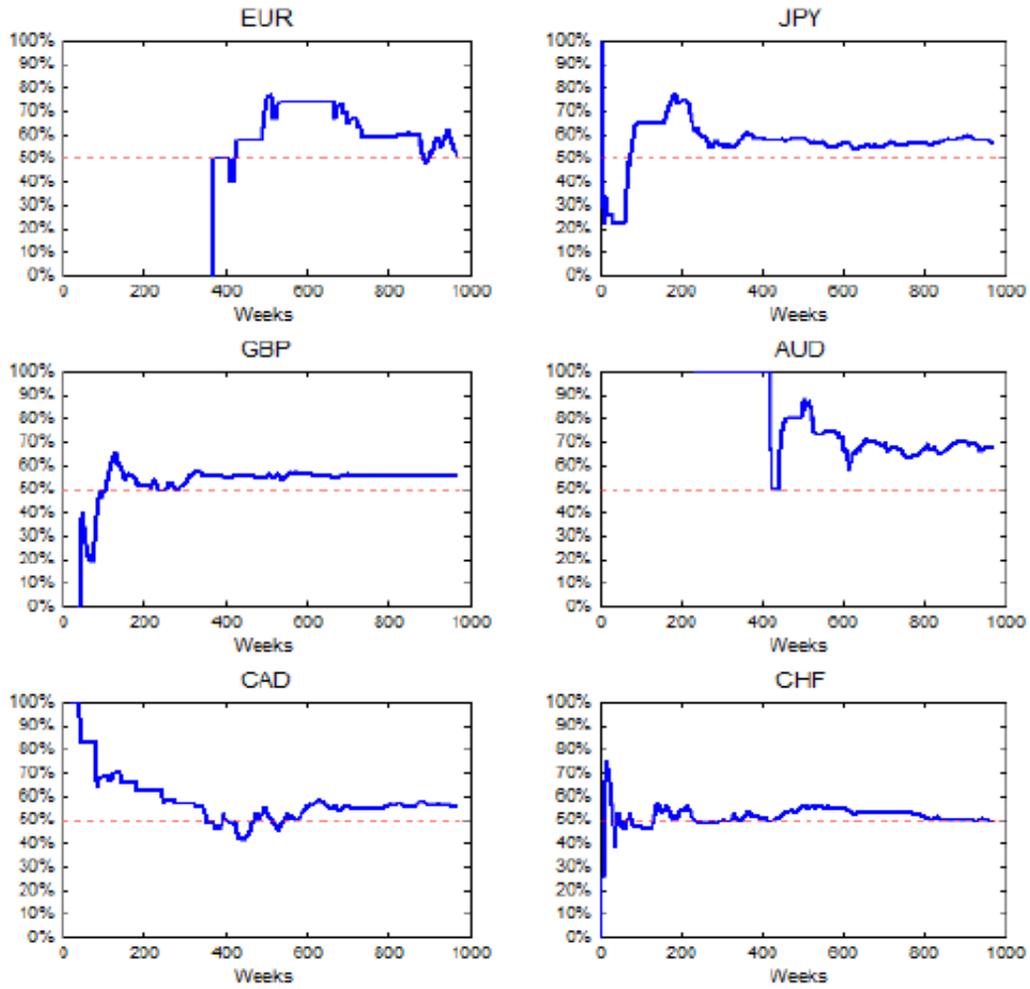
Notes: 1. This graph presents the estimated AR(1) coefficient of the auto-regressive MSM. The red solid line depicts the estimators. 2. The information on the rolling window sizes in the estimation of the MSM are described in the note of Figure D.1.

Figure D.3. Evolution of Cumulative Forecast Success Ratio (h=1m)



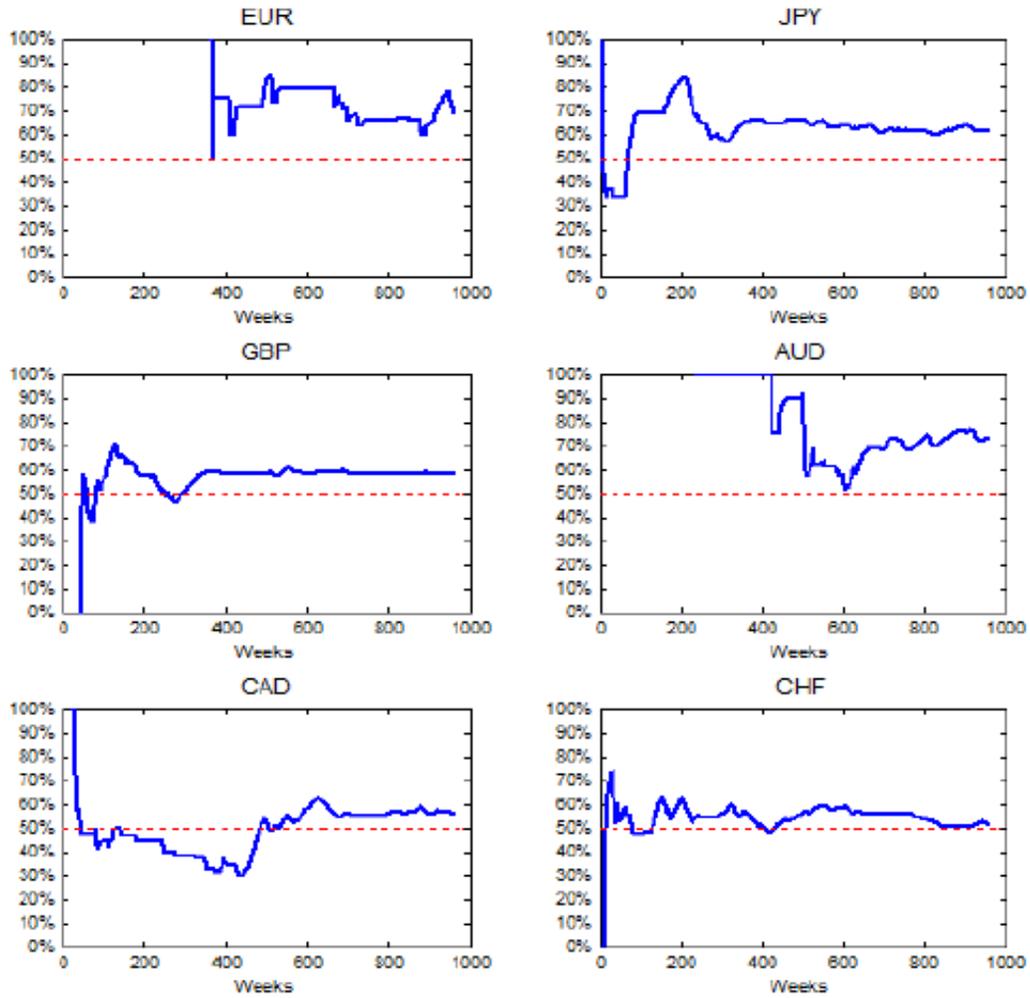
Notes: 1. $t=0$ is the first week in which h month ahead forecasts are generated. 2. The forecast success ratio is defined the number of successful depreciation or appreciation forecasts divided by the total number of depreciation and appreciation forecasts. When we predict no change $D_{t,h}=0$, we do not count it in the calculation of forecast success ratio.

Figure D.4. Evolution of Cumulative Forecast Success Ratio ($h=3m$)



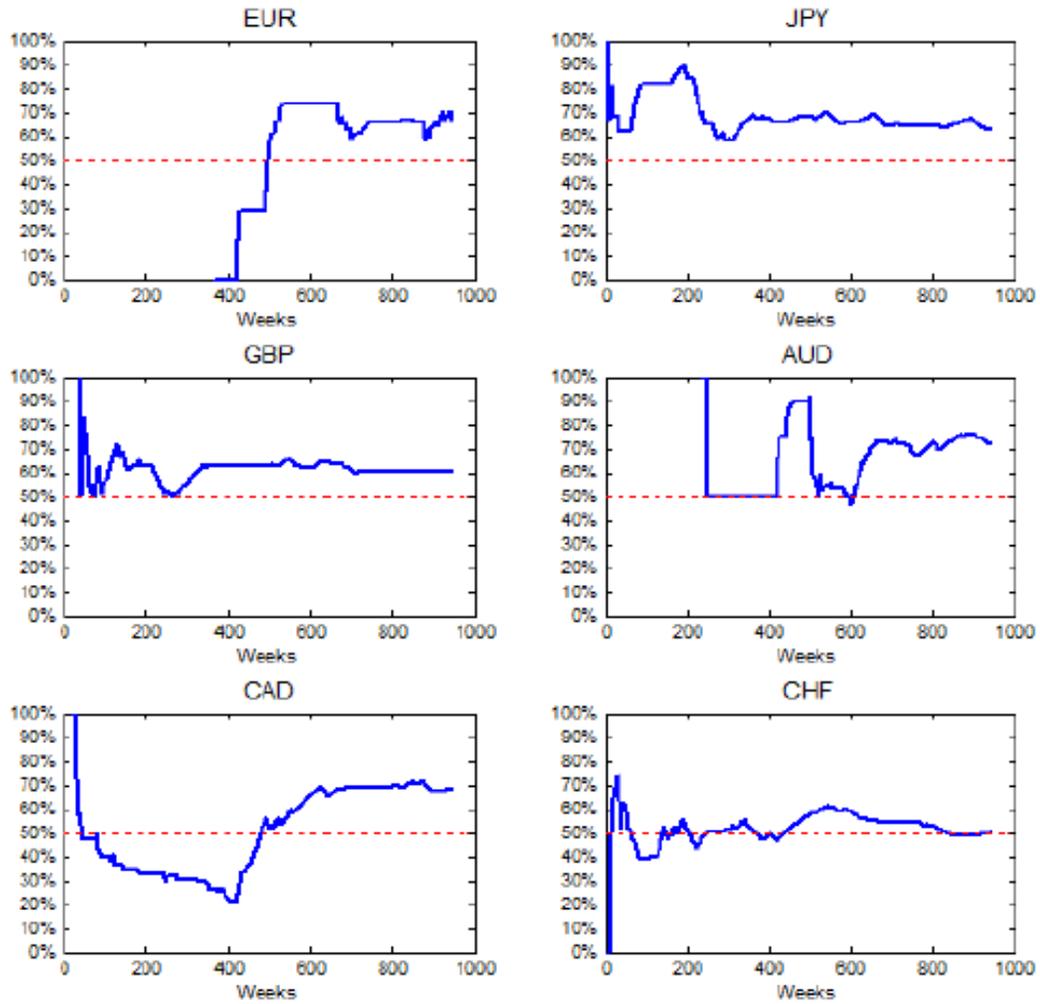
Notes: The cumulative forecast success ratio is described in the note of Figure D.3.

Figure D.5. Evolution of Cumulative Forecast Success Ratio (h=6m)



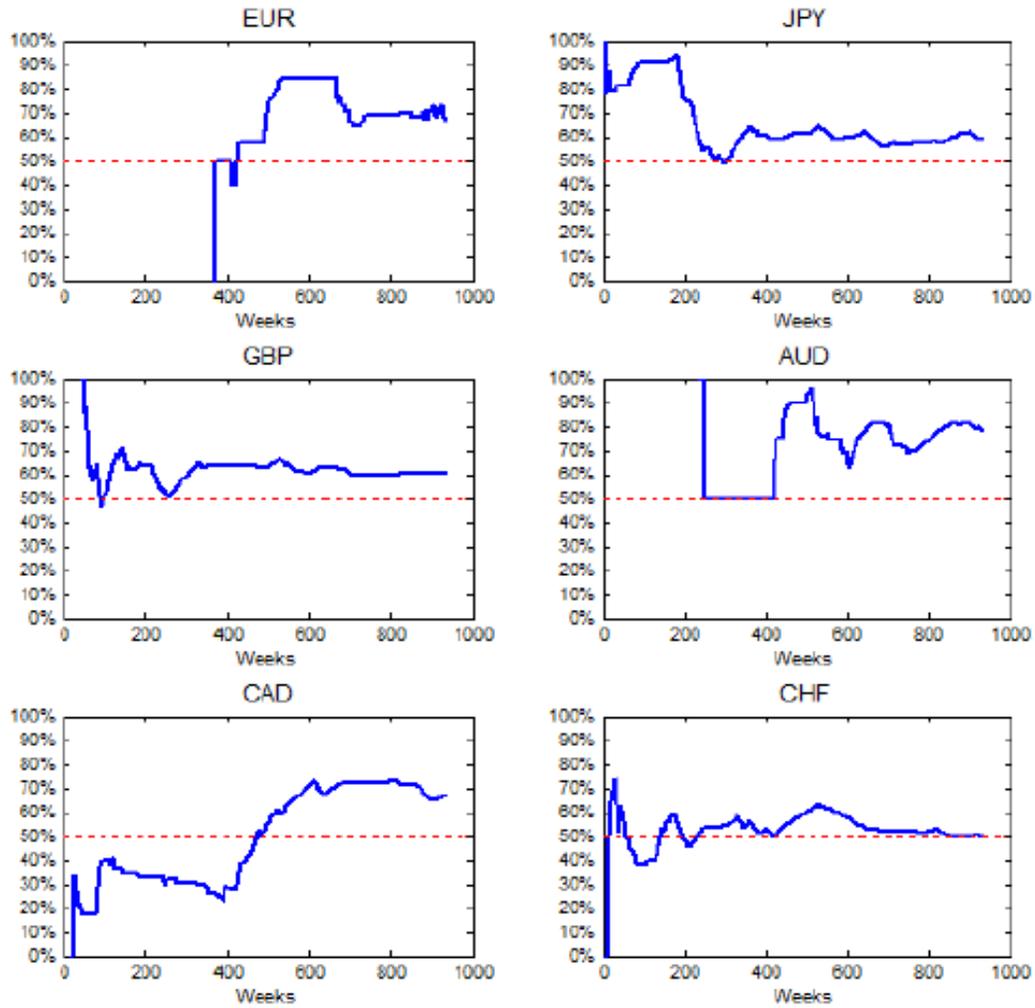
Notes: The cumulative forecast success ratio is described in the note of Figure D.3.

Figure D.6. Evolution of Cumulative Forecast Success Ratio (h=9m)



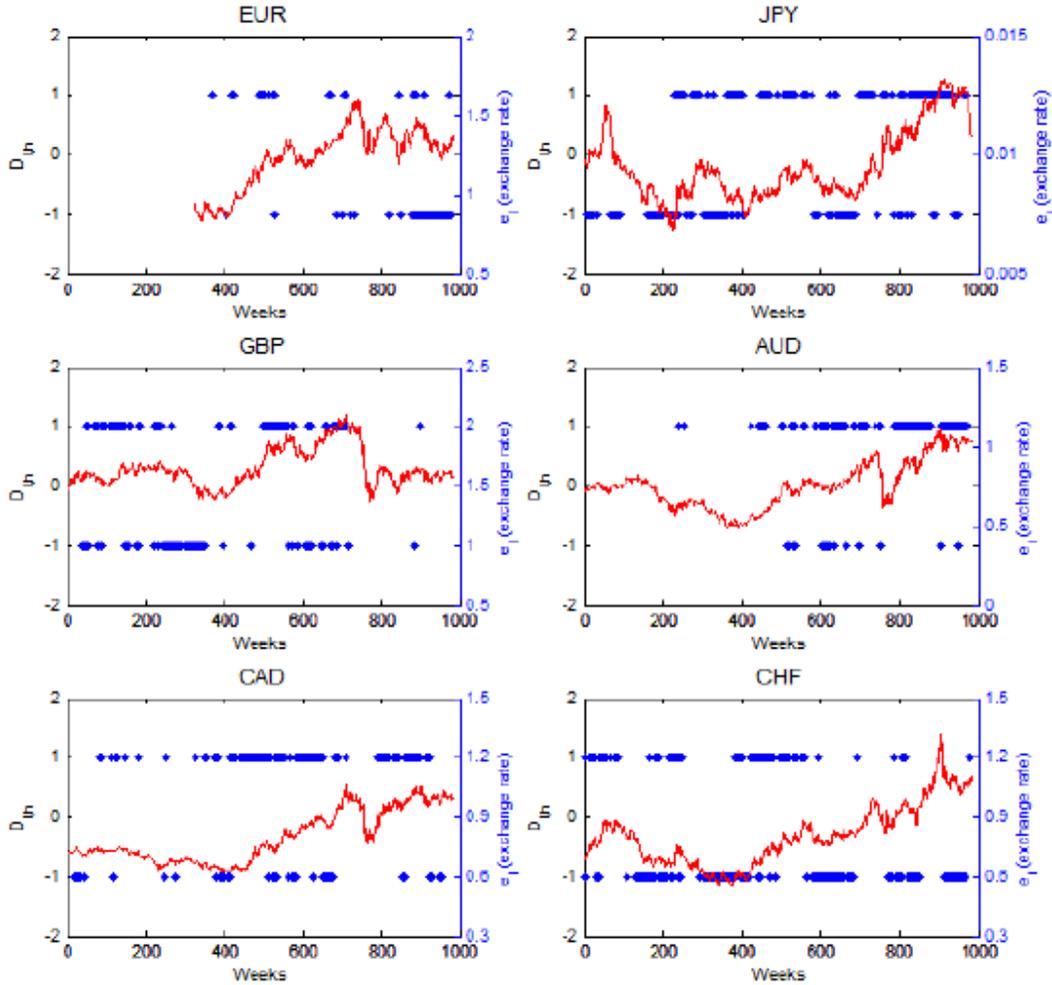
Notes: The cumulative forecast success ratio is described in the note of Figure D.3.

Figure D.7. Evolution of Cumulative Forecast Success Ratio (h=12m)



Notes: The cumulative forecast success ratio is described in the note of Figure D.3.

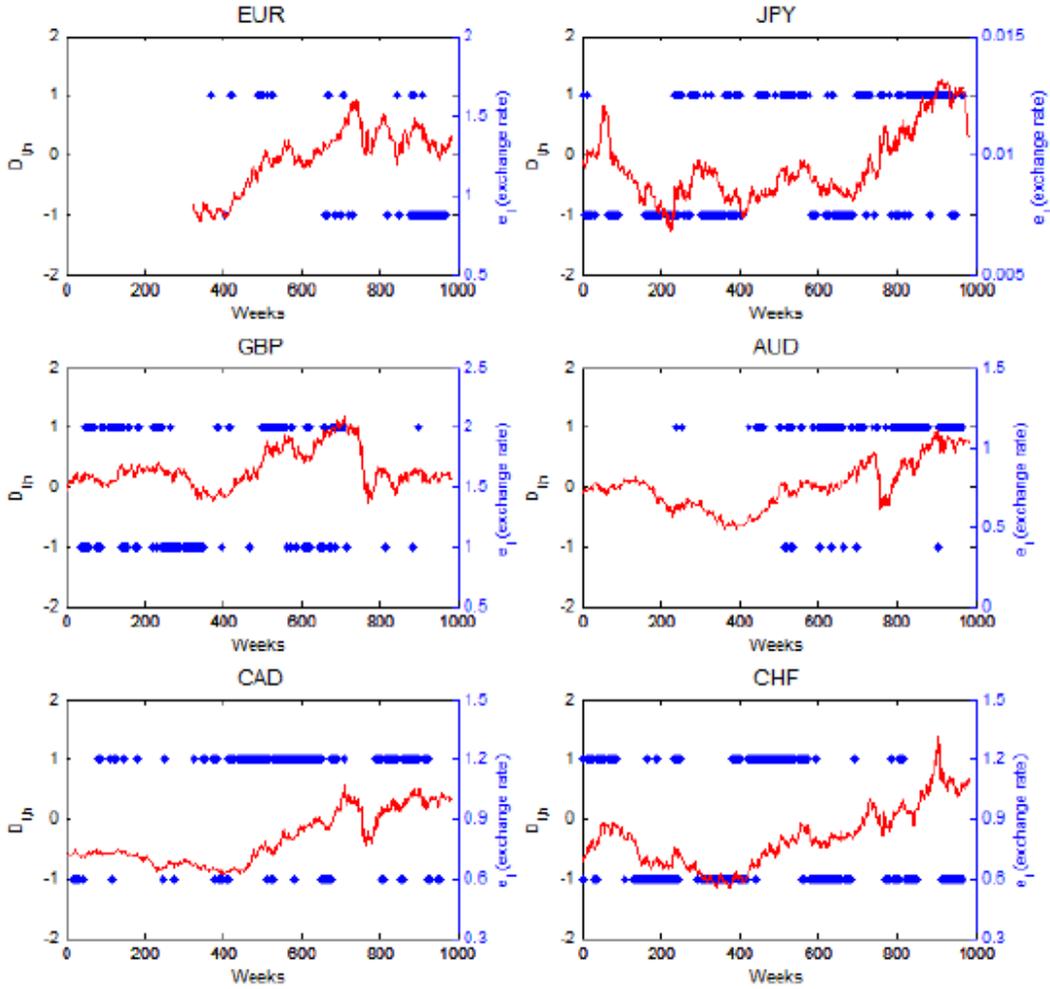
Figure D.8. Evolution of $D_{t,h}^{(i)}$ ($h=1m$)



Notes: 1. This graph presents the $D_{t,h}^{(i)}$ on which our directional forecasts are based. $D_{t,h}^{(i)}$ (blue scatter plot) is on the left axis and the exchange rate (red line) on the right axis. 2.

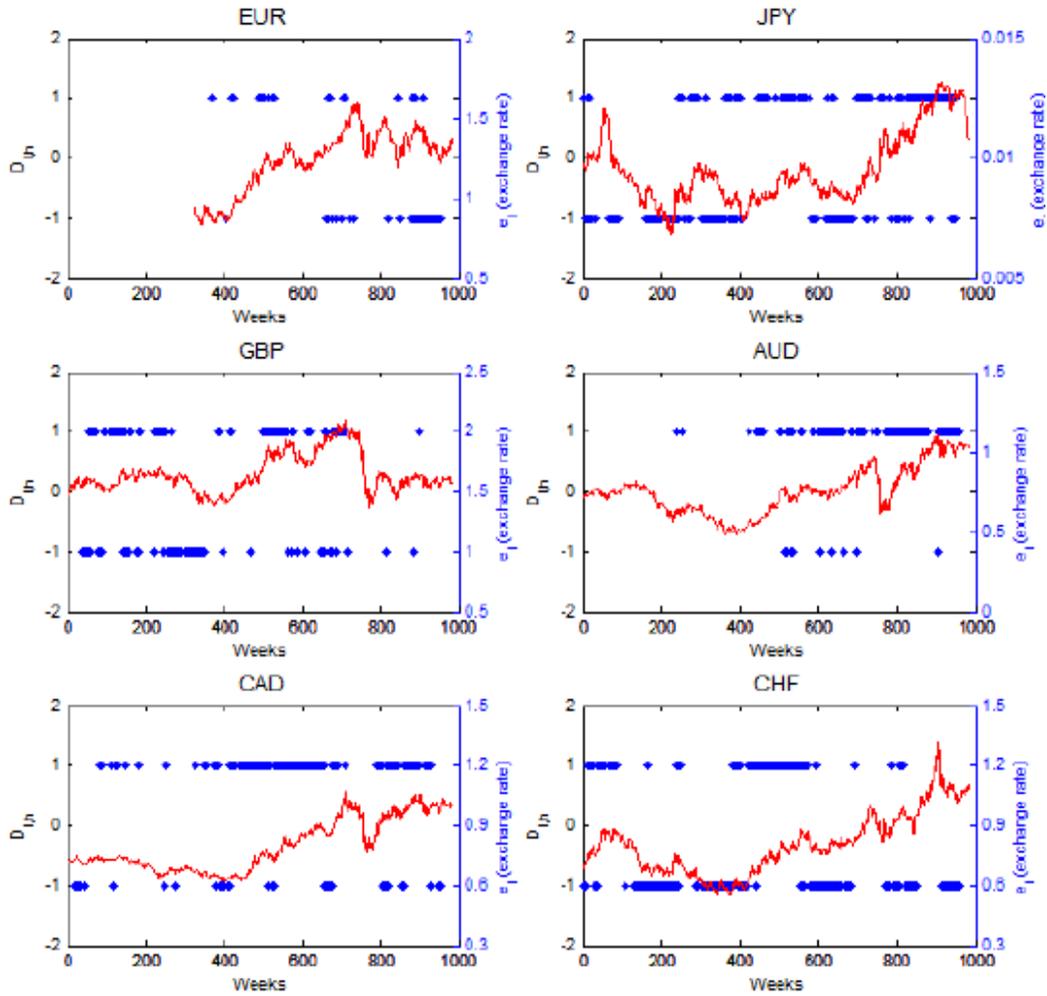
Blue scatter plots can take values -1 or 1. No scatter plots on the given weeks indicate $D_{t,h}^{(i)}=0$. -1 (or 1) predicts depreciation (or appreciation) of the given currency against US Dollar over the h month forecasting horizon. No scatter plots on the given weeks mean that our directional forecasts predict no change over the same forecasting horizon.

Figure D.9. Evolution of $D_{t,h}^{(i)}$ (h=3m)



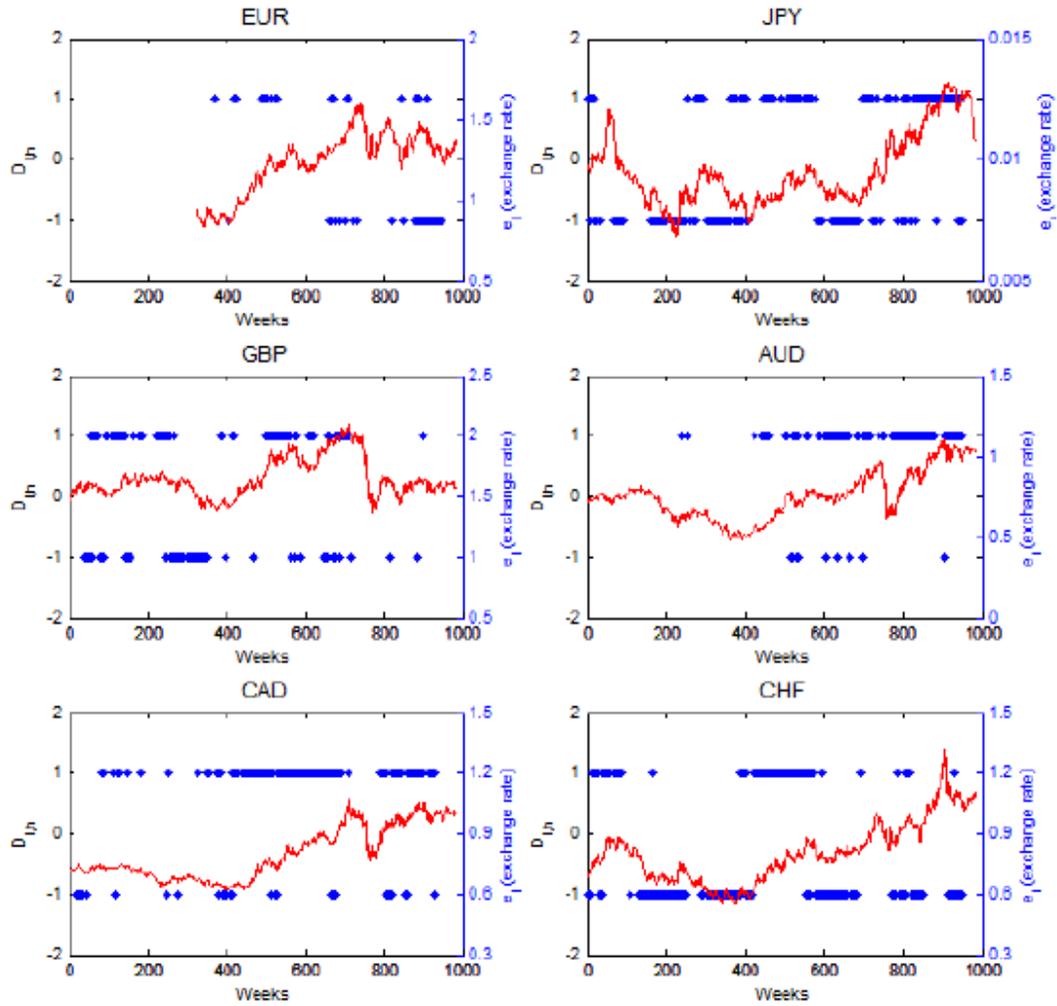
Notes: The information on the graph is described in the note of Figure D.8.

Figure D.10. Evolution of $D_{t,h}^{(i)}$ (h=6m)



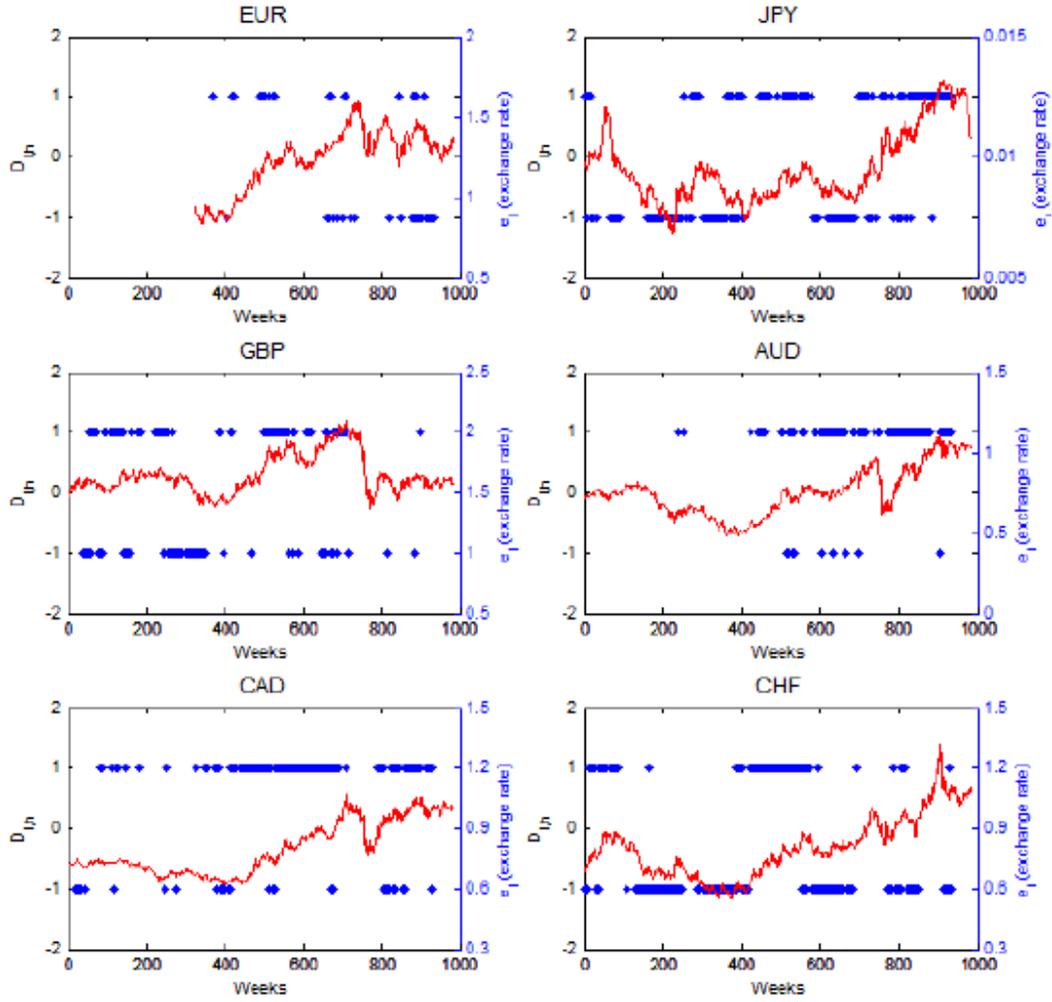
Notes: The information on the graph is described in the note of Figure D.8.

Figure D.11. Evolution of $D_{t,h}^{(i)}$ (h=9m)



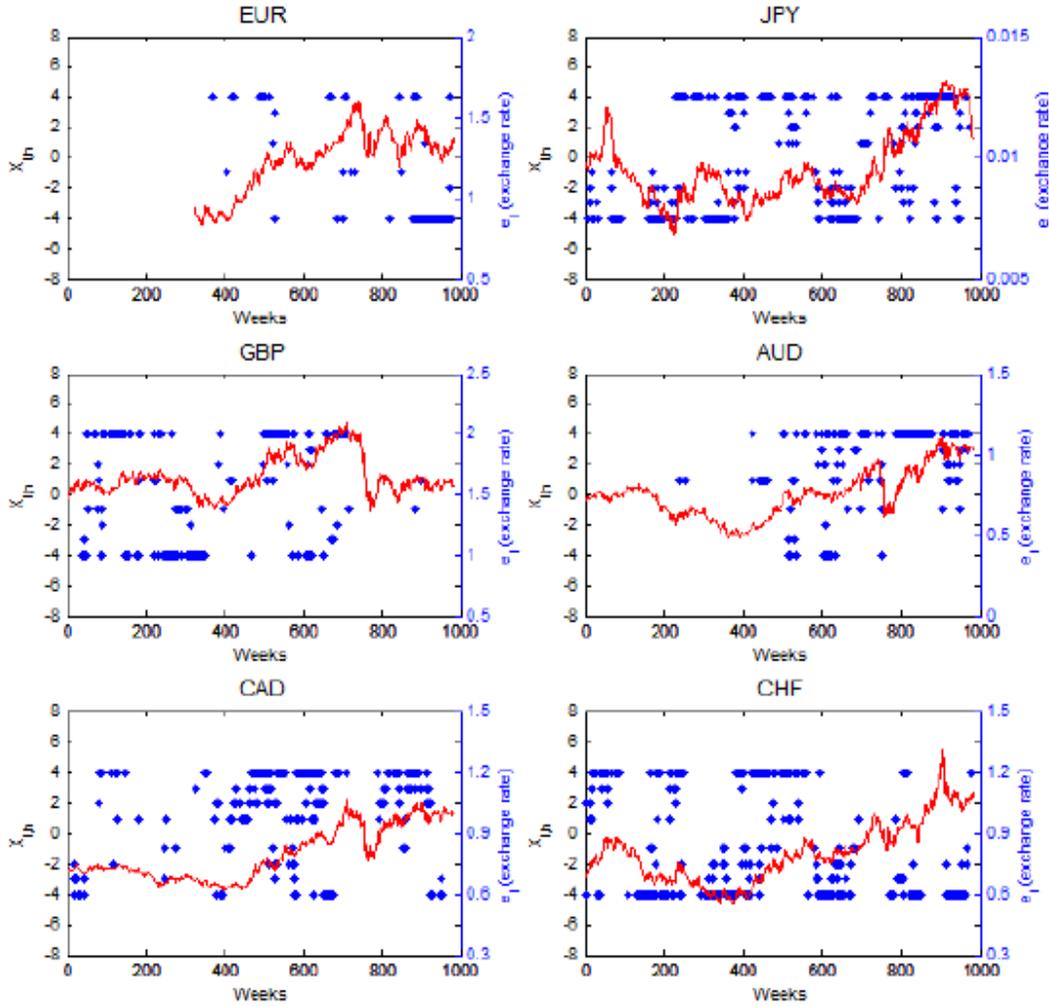
Notes: The information on the graph is described in the note of Figure D.8.

Figure D.12. Evolution of $D_{t,h}^{(i)}$ (h=12m)



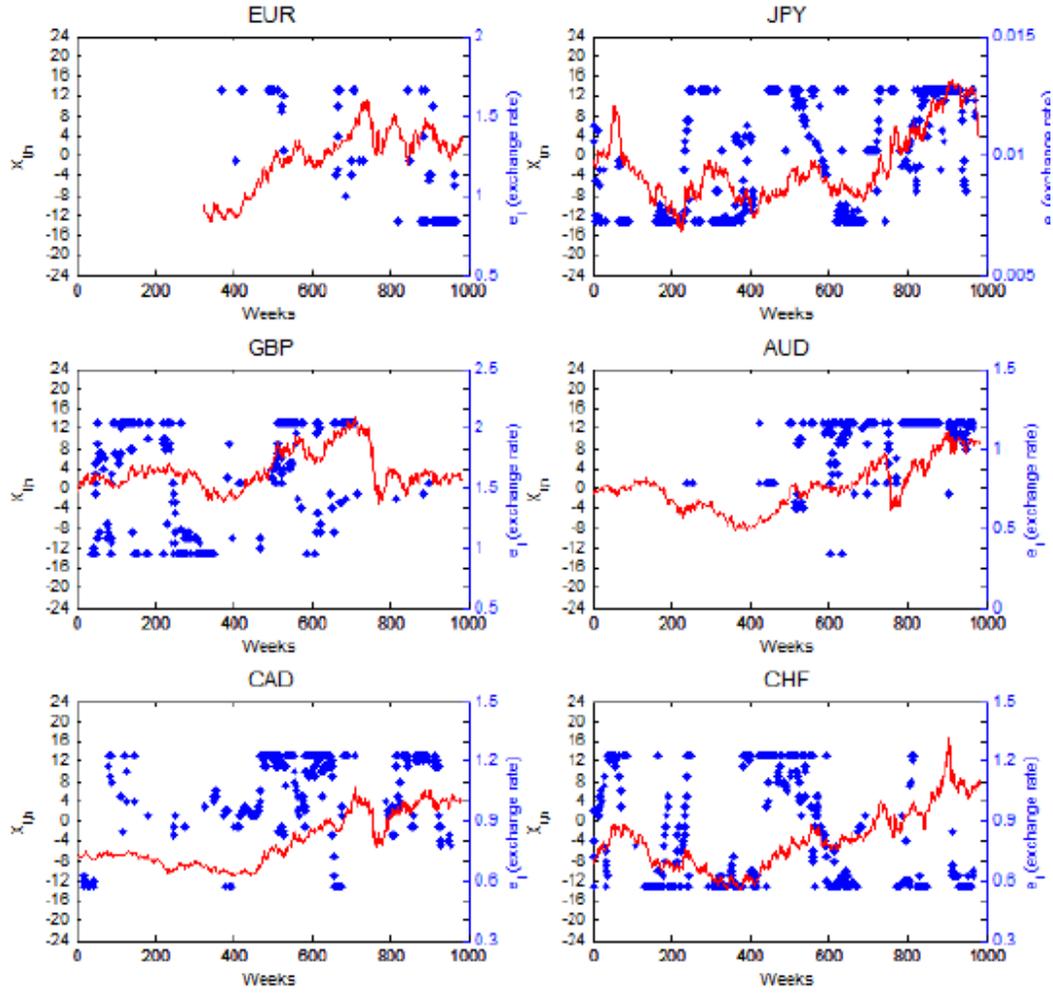
Notes: The information on the graph is described in the note of Figure D.8.

Figure D.13. Evolution of $X_{t,h}^{(i)}$ ($h=1m$)



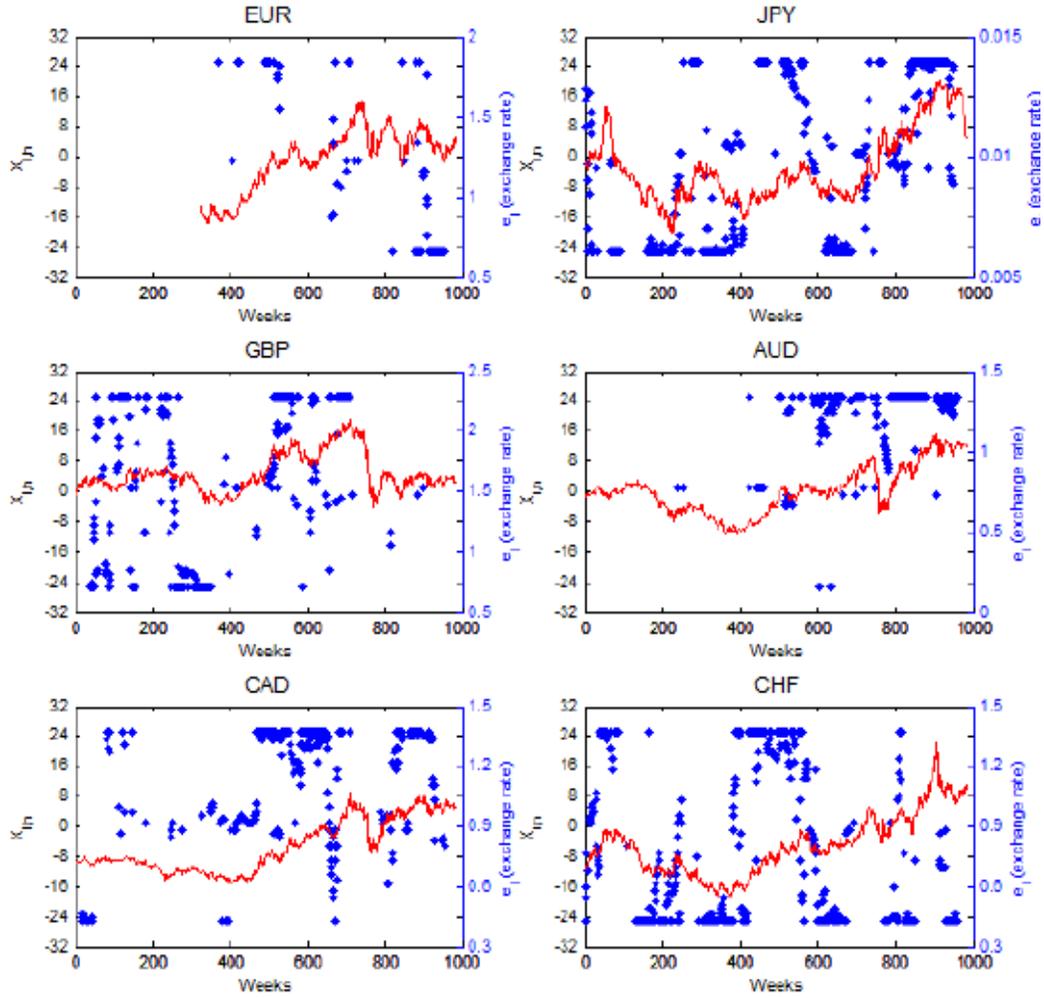
Notes: 1. This graph presents the $X_{t,h}^{(i)}$, which is the net number of periods with predicted speculators' accumulation (decumulation) over the following h months. $X_{t,h}^{(i)}$ (blue scatter plot) is on the left axis and the exchange rate (red line) on the right axis. 2. Blue scatter plots can take values from $-h$ to h except for 0. No scatter plots on the given weeks indicates that $X_{t,h}^{(i)}=0$. $-h$ (or h) indicates the magnitude of depreciation (or appreciation) over the h forecasting horizon. No scatter plots on the given weeks means that our directional forecasts predict no change over the same forecasting horizon.

Figure D.14. Evolution of $X_{t,h}^{(i)}$ (h=3m)



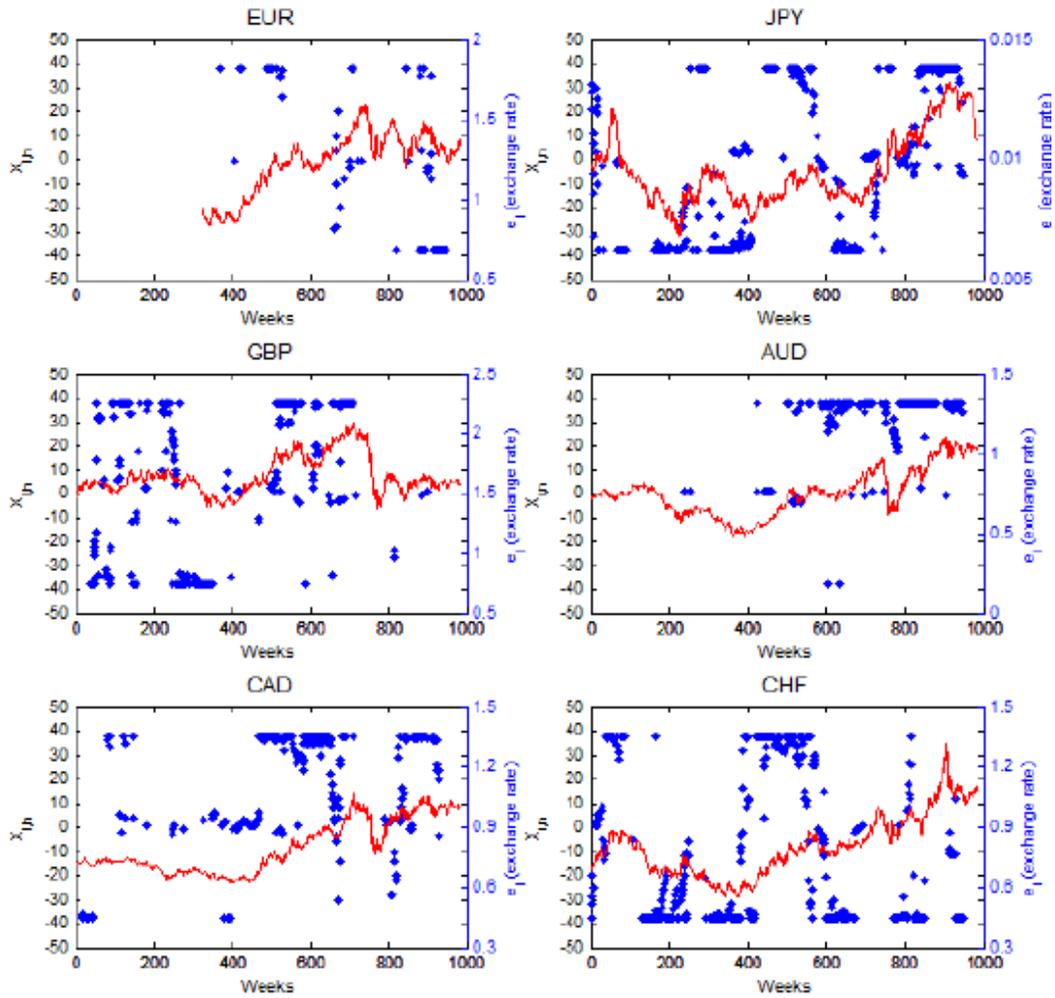
Notes: 1. The information on the graph is described in the notes to Figure D.13.

Figure D.15. Evolution of $X_{t,h}^{(i)}$ (h=6m)



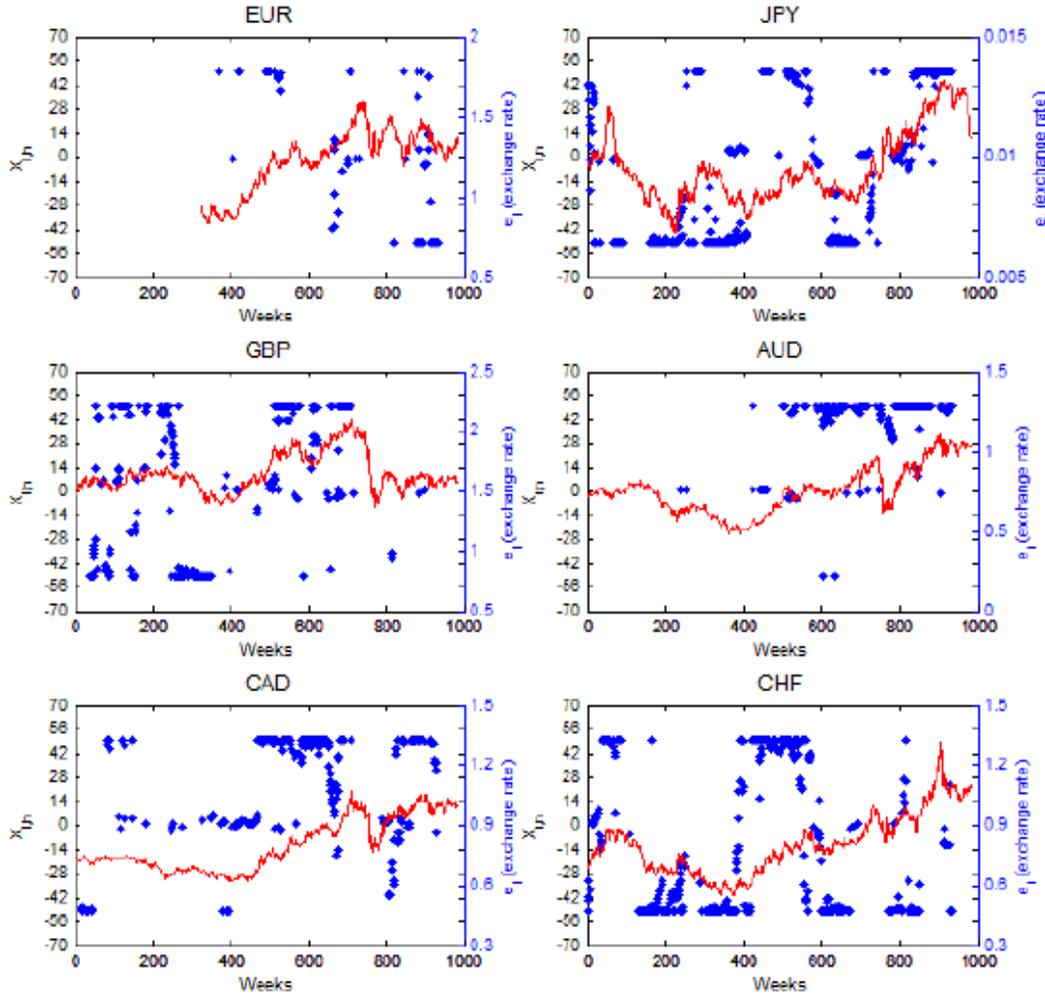
Notes: 1. The information on the graph is described in the notes to Figure D.13.

Figure D.16. Evolution of $X_{t,h}^{(i)}$ (h=9m)



Notes: 1. The information on the graph is described in the notes to Figure D.13.

Figure D.17. Evolution of $X_{t,h}^{(i)}$ (h=12m)



Notes: 1. The information on the graph is described in the notes to Figure D.13.